ircam
Centre Pompidou

Software Documentation

AudioSculpt

User's Manual

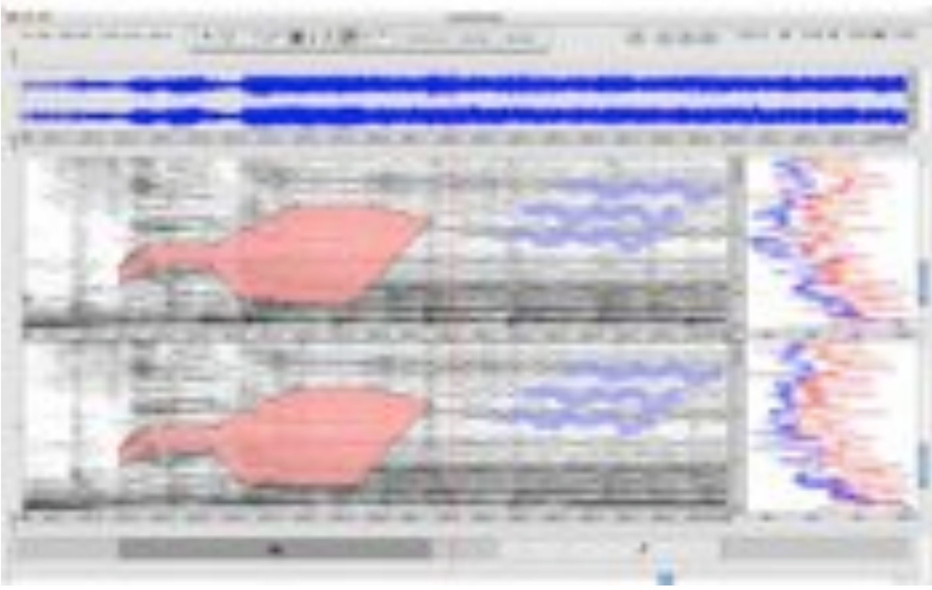

Alain LITHAUD

AudioSculpt: User's Manual
by Alain LITHAUD

Edited by Karim Haddad

Published Date: 2003/10/07 08:49:59

This manual was produced under the editorial responsibility of Karim Haddad, Marketing and Communication Dept. , IRCAM.
This manual was translated by Justice Olsson

Conception and developpement

AudioSculpt:
Niels Bogaards. Contributions : Philippe Depalle, Nicholas Ellis, Peter Hanappe, Alain Lithaud, Marc Locascio, David Ralley, Alberto Ricci, Chris Rogers, Xavier Rodet, Hans Tutschku

SuperVP:
Axel Roebel. Contributions : Philippe Depalle, Gilles Poirot.

PM2:
Axel Roebel.

libsndfile:
Erik de Castro Lopo

EASDIF/SDIF:
Diemo Schwarz, Patrice Tisserand, Fabien Tisserand, Xavier Rodet, Axel Roebel

This documentation corresponds to version 2.3.2 or higher of Audiosculpt.

Marketing and Communication Dept.,
IRCAM
1 place Igor-Stravinsky

F-75004,Paris
FRANCE
Phone:(33) (1) 44 78 49 62
Fax: (33) (1) 42 77 29 47
Email: palumbo@ircam.fr

Send comments or suggestions to the editor:

Karim Haddad
Marketing and Communication Dept.,
IRCAM
1 place Igor-Stravinsky
F-75004,Paris
FRANCE
Phone:(33) (1) 44 78 16 05
Email: haddad@ircam.fr

For more information:
http://www.forumnet.ircam.fr

Table of Contents

1 Introduction

AudioSculpt is a powerful sound analysis and processing tool. AudioSculpt allows various analysis types of a sound, and can store the results SDIF format files (or in text files, if required) which can be exploited by other software. The SDIF format can be converted into text (or, conversely, from text) by means of the "Droplets" that come with AudioSculpt.

Certain analysis results can be displayed via the spectrogram. AudioSculpt allows for the defining and the application of various sound processes: various types of filtering, stretching, transposing etc. It also allows the user to manage these treatments via a treatment sequencer.

AudioSculpt can work with mono sounds as well as multichannel sounds. It accepts the AIFF, SDII and WAV formats, in 8, 16,24,32 bits, and can handle frequencies up to 192 kHz. It can also handle files by "floating" AIFC and WAV in 32 bits.

2 Hardware requirements

AudioSculpt 2.3 was developed for Macintosh. Its minimum requirement is the Mac OS X. 10.3 system (Panther).

Processing speed and the ability to process in real-time depend on the CPU and its speed, as well as the amount of RAM.

Required disk space naturally varies according to the work carried out.

A reminder: a file containing one minute of stereo sound (16 bits at 44 100 Hz) is around 10 megabytes.

It is also necessary to take into account temporary files size: a FFT file will take up about twice as much space as the sound file, using typical settings.

AudioSculpt functions with "CoreAudio" and can therefore uses various sound cards for sound output.

The current version is 2.3.2. The last compatible version with Mac OS nine and Mac OS X. is AudioSculpt 2.0.4b.

> **Note:** you can discuss AudioSculpt and its applications by enrolling yourself in the user's list (both in French and in English) at the following address:
>
> http://listes.ircam.fr/wws/lists/english_private

3 Installation

All you have to do is mount the disk image (.dmg) and copy the dossier "AudioSculpt 2.3" to the location desired on your hard disk. This dossier contains the following elements:

- **AudioSculpt 2.3.app** : the application itself.
- **Documentation** : self explanatory.
- **Droplets** : a folder containing various droplets (and the corresponding kernels) which among other things allow for SDIF to Text conversions and vice versa.
- **Fft** : an empty folder for storing Fft files (temporary or otherwise).
- **Fundamental** : an empty folder for storing files (temporary or otherwise).
- **Kernel** : a folder containing the application's kernels (SuperVP, Super Phase Vocoder, pm2 and the additive synthesis motor) as well as the documents necessary for operation of the application.
- **Markers** : an empty folder for storing marker files.
- **Sounds** : a folder containing the sound file "Africa.aiff" by way of an example. Used for storing sounds.
- **SpectralEstimates** : an empty folder for storing SDIF spectral estimation files.
- **Temp** : an empty folder for storing temporary files.
- **Treatments** : an empty folder for storing processing files.

It is extremely important not to move any of the above their spheres or documents except for "Documentation" and "Droplets", or AudioSculpt will not run properly.

When first launched, AudioSculpt opened the window inviting you to type in your name and your code: please follow the instructions that come with the letter accompanying the CD-ROM very carefully.

> **Note:** in the event of a problem, and if an earlier version of AudioSculpt was installed, the preferences file "**AudioSculpt 2 Prefs**" must be discarded. It is in the "**Preferences**" folder contained in the "**Library**" folder in your account (your house).

> **Important:** file names (whether sound or other files) must not contain "/" (forward slash) or any other diacritic characters.

4 Uninstalling

Uninstalling is extremely simple: just drag the folder containing the application to the trash can and empty it.

It can also if desired discarded the "**AudioSculpt 2 Prefs**" file in the "**Preferences**" folder, contained in the "**Library**" folder in your account (your house).

5 Getting Started

5.1 Launching AudioSculpt

AudioSculpt is started in the usual way: either by double clicking the icon, or by dragging a sound file onto it. The first method (clicking on the icon) will open those windows that were ticked in the "Window" menu ("Tools", "Inspector", "Sonogram Display", "Grid Settings"). Please see section 31.6

Note: the " tabulator" key tunnels between masking and displaying those palette windows.

Then, go to the File menu and choose "Open. . . ".

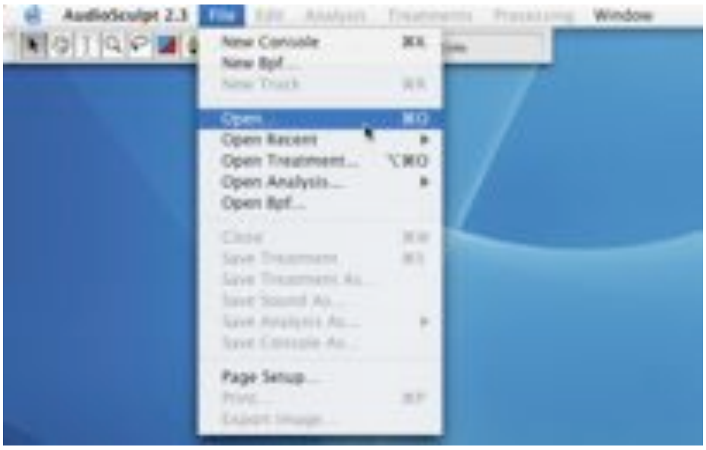

In the standard dialog box, choose a sound.

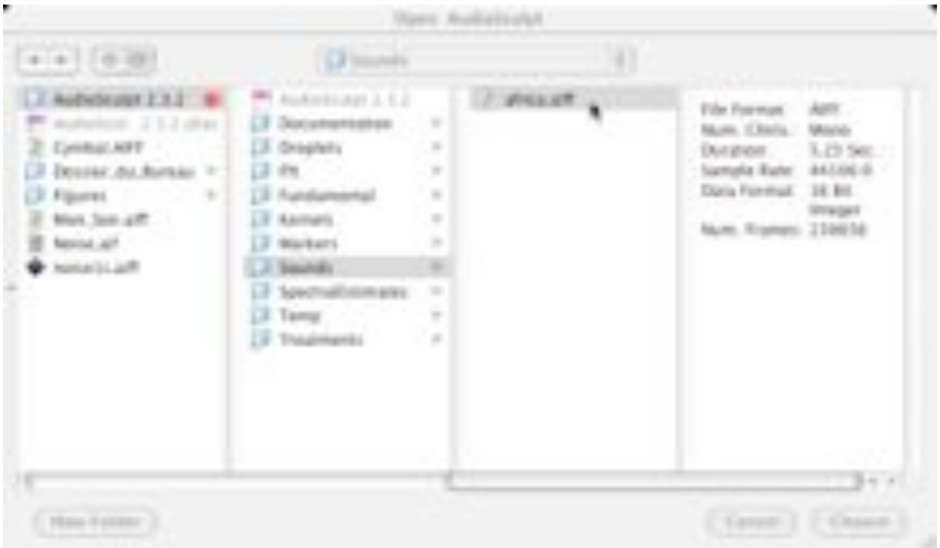

This is what the main AudioSculpt window looks like, carrying the name of the sound:

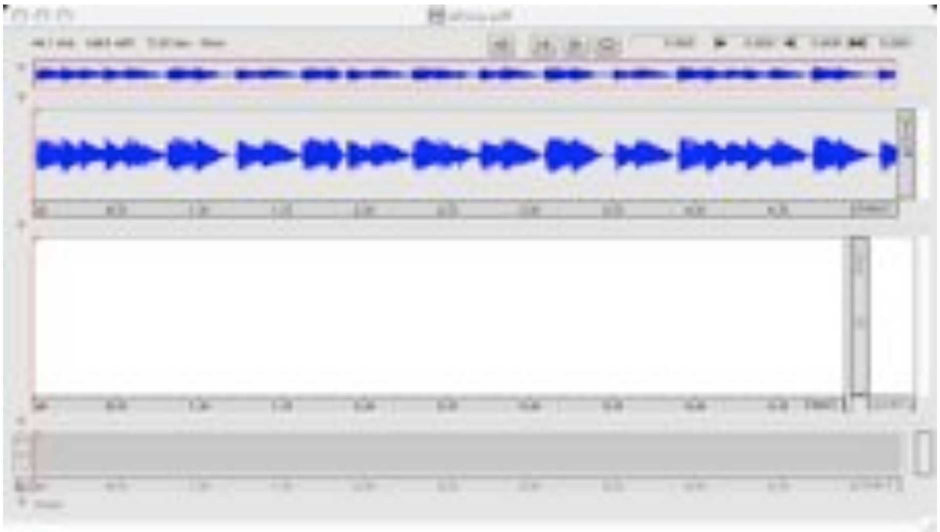

This window is directly opened by dragging a sound onto the application.

You can open a recent sound file by choosing it in the submenu via "Open recent" in the "File" menu.

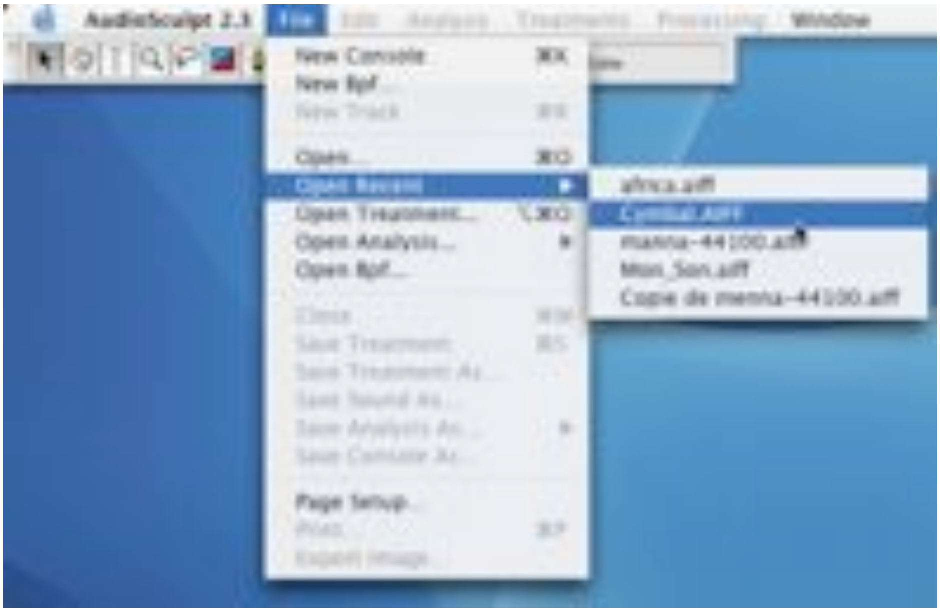

5.2 The AudioSculpt window

This window is resized in the usual manner (the square, bottom right).

At top left information is displayed about the currently open sound file (sampling rate, quantisation resolution, formats, duration, number of channels). On the right are four buttons for

playing back the sound, followed by four fields that can be edited. They are for setting the position of the cursor and for selecting it. Please consult sections 6, 14 et 15.

Zone 1 (upper zone) shows the entire sound (as a wave form). If the sound is multichannel the sum of the channels is displayed. The red frame represents the visible part of the sound in zone 2.

In order to play back the sound, press the spacebar. To stop play back of the sound, press it again. You can also use the third button at the top of the window. For more detail, please consult section 6.

Zone 2 is the "**window into the sound**": It allows you to navigate around the sound by zooming in and out. Note that initially it shows the entire sound.

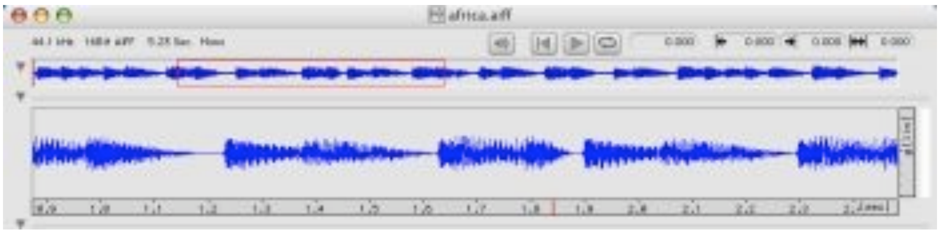

Zone 3 displays the sonogram as well as the spectrum, on the right.

La zone 3 est l'endroit où s'afficheront le sonagramme ainsi que le spectre, à droite.

Zone 4 is the sequencer : this is a set of tracks, where processing takes place.

Zone 5 is initially folded away (please see below) and indicates "Ready" : this is the SuperVP console. This is where command lines and comments generated by SuperVP are displayed. This console is linked to the open sound file.

Note: zones 2, 3 all have identical time axes. They are synchronized and you can navigate around them using the horizontal "elevator" at the bottom of the window. It is displayed when you zoom inside the sound.

Note: All zones have a small triangle, allowing you to fold/unfold them.

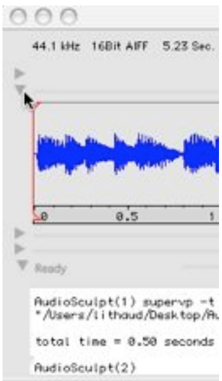

Zone 1, 2 and 3 can be vertically resized: move the pointer over the bar between the two zones.

Then click the mouse and the pointer will change, allowing you to move the bar up or down.

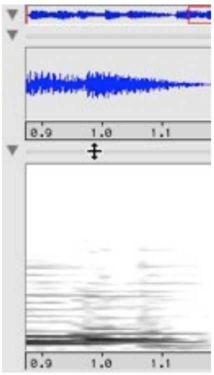

There is a similar function allowing you to set the respective widths for the sonogram and the spectrum.

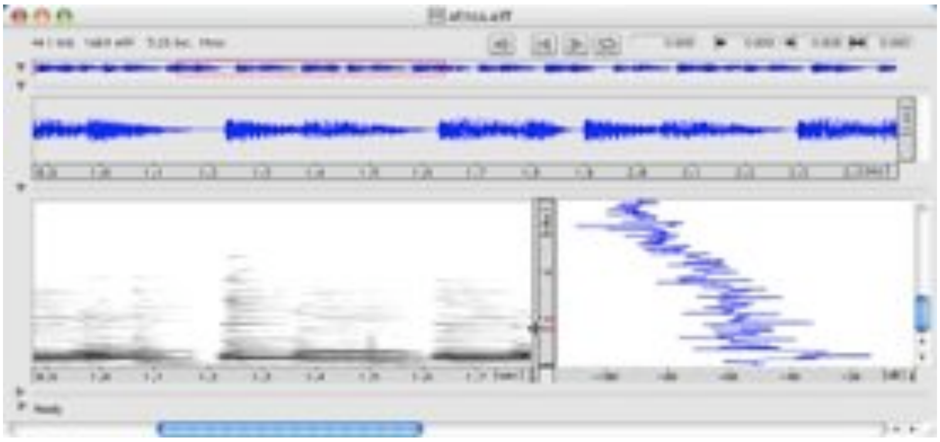

Note: if you stretch the sonogram all the way to the right, you will only see the bar in the timescale, under the frequency scale.

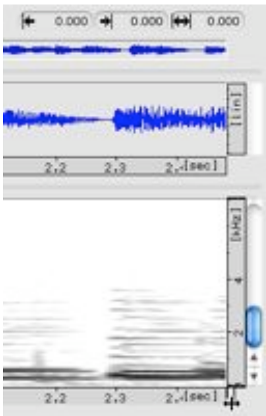

6 Sound playback

To play back the sound, press the space bar and to stop, press it again.

Cursor (also known as song position line) movement is shown as a red line over all the zones, and the position is indicated in the field at the top of the window. Please consult section 14.

If part of the sound is selected, then only that part will be played back. There are three buttons used for play back at the top of the window:

Clicking on the first button will place the cursor at the start of the sound or the start of the selection. A second click will bring it back to the start of the sound (if there is a selection) but only the selection will be played back (keyboard shortcut: "**Return**").

The second button: "play/stop" fulfills the same function as the space bar. The third button allows you to play back the selection in a loop. If there is no selection, the entire sound will be looped.

In order to play back the sound from a given point, cancel the selection by clicking on the sonogram with the cursor tool. Please consult sections 11.1 and 15). The sound will play back from the new cursor position. To play back the sound from its beginning, just move the cursor to the beginning by means of the first button described above.

The cursor may be placed at any desired position (please consult section 14).

When you place any tool (except for the hand tool) in zone 1 (entire sound) it will change into a loudspeaker with a vertical line (please consult section 11.2.2). Thus it is possible, by clicking, to play back from any point whatsoever.

If you have zoomed forward into the sound, the "follow playback" item in the "Window" menu will cause the sound to scroll along with the cursor movement. If this item is unchecked, the "window in the sound" remains immobile.

7 Sound outputs

7.1 Integrated outputs

AudioSculpt includes support for the Macintosh integrated soundcard outputs. If you use them:

- mono sounds will be played in mono through both the Macintosh sound outputs;
- stereo sounds (or two channel sounds) will be played in stereo through the two outputs (left and right) of the MacIntosh;
- for multichannel sounds (other than stereo or two channel) only the first two channels will be played, in stereo, through the two Macintosh outputs (left and right).

7.2 External soundcards

AudioSculpt functions with "CoreAudio" and can therefore use various sound cards.

The Apple "audio MIDI setup" (in the "utilities" folder) must be configured in order to use external sound cards.

> **Note:** The external sound card driver may require additional configuration.
>
> For example, in the case of the "Digi 002" (Digidesign) AND version 6.4 of Protools LE, you have to open the "Digidesign CoreAudio setup" utility (it is in the "Digidesign" folder) and add the "AudioSculpt" application to "Supported Applications. . . ". Note that "Digidesign CoreAudio setup" does not recognize super VP as an application. Therefore the Macintosh integrated outputs will run in "real-time mode".

8 Sound files

AudioSculpt supports mono sounds as well as multichannel sounds. It will accept the AIFF, SDII and WAV formats, in 8,16,24, the 22 bits. It recognizes sample rates all the way up two 192 kHz.

AudioSculpt can work with AIFC and WAV "floating" sound files in 32 bit architecture.

9 Display scales

You can choose certain display scales in the AudioSculpt or the Bpf Windows. By control clicking on a scale, you open a pop-up menu that will allow you to choose.

- **Time scales** for the sonogram, for the sequencer tracks as well as for the Bpf's condescending transposition, stretching/compression or gain:

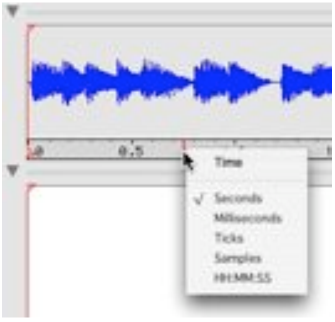

when you choose a time scale for a zone in the AudioSculpt window, it applies to all zones. By contrast, the time scale that you choose for a Bpf only applies to the Bpf.

- **Frequency scales** for the sonogram and the spectrum together, as well as for the "breakpoint filter" Bpf:

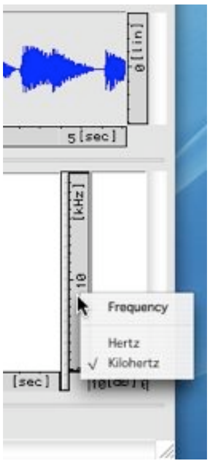

- **Transposition scale** (transposition via Bpf) :

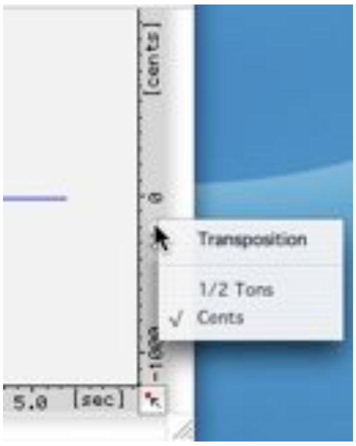

- **Amplitudes scales** : linear for the " window into the sound" and logarithmic (in dB) for the spectrum as well as for the " breakpoint filter" and "gain" Bpf's. This is not subject to choice.

10 Zooming

These manipulations apply to all scales (horizontal and vertical scales in the AudioSculpt and Bpf windows).

10.1 Managing the zoom scales

To magnify the scale around any value, simply hold down the mouse button while dragging the mouse to the right. To reduce the scale, drag the mouse to the left.

Use the same maneuver but in addition hold down the "**Apple**" key (" **command**") in order to vary the zoom from the leftmost value in the window.

By placing any of the tools on any value in any of the scales, and holding down the mouse button, the scale can be made to stretch out (increase) around that value for as long as the mouse button is held down.

Double click on the ruler to return to the initial size.

10.2 The magnifying tool

The magnifying tool (please consult section 11.6) allows you to select part of the sonogram (and/or all of the spectrum) or of the sound ("window into the sound") and to resize it up to the size of the window. Double click on the ruler to return to the initial size.

The magnifying tool can also become negative: press the "**option**" ("**alt**"): zooming will now be zooming out. Double click on a scale to return to the initial size.

Place the magnifying tool on the " window into the sound" or on the sonogram (or the spectrum) while "**option clicking**" (**alt clicking**) to make the magnifying tool negative. It will now zoom out in jumps from the center of its position. As soon as you zoom in the sound or the sonogram, the red rectangle in the upper zone shows the borders of the visible sound in zone 2.

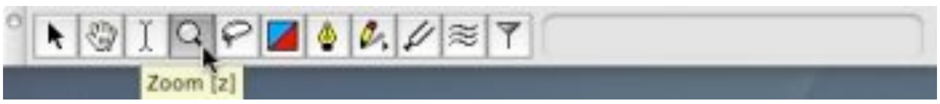

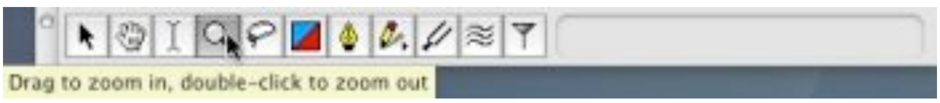

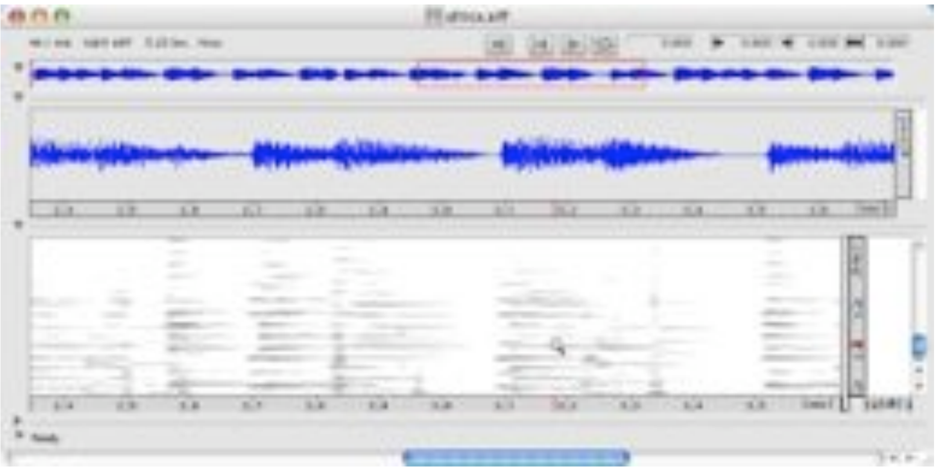

To navigate around the sound and the sonogram place the hand tool (please see section 11.4) on the red rectangle and hold down the mouse button.

You can move the sonogram image in any generation in its window using the hand tool (the "window into the sound" display will follow accordingly).

The hand tool also allows you to grab the Bpf image and move its in any direction in its window.

In the "Window" menu, "the Optimize Sonogram" item allows you to display the sonogram and the "window into the sound" in full screen mode (zones 2 and 3), or to return to the initial size. The keyboard shortcut for this is "**Apple +U**" (**command +U**").

11 The Toolbox

11.1 The toolbox

The toolbox is displayed if you checked "Show Tools" in the "window" menu (it is best to check this option permanently).

The" **tab**" key will toggle this toolbox between hide/show, as well as the three others ("Inspector", " Sonogram Display", " Grid Settings") on condition they have been checked.

The pointer position is displayed on the right, in the tool bar. It can be "fixed" by pressing the "f" key, which allows it to be copied and pasted elsewhere (please see section 11.2.1).

Keyboard shortcuts make it easy to rapidly move from one tool to another:

- **a** = **A**rrow
- **h** = (**H**and)
- **b** = (**B**eam - for selection)
- **z** = (**Z**oom)
- **l** = (**L**asso)
- **s** = (**S**urface)
- **r** = (**R**egion)
- **e** = (p**E**ncil)
- **d** = (**D**iapason)
- **c** = (harmoni**C**)
- **m** = (**M**arker)

Hover over a tool to display its name and keyboard shortcut (these are the "help tags" found in many applications). Press the "**Apple**" key (**command**) for a more detailed description where available.

11.2 Features common to all tools

11.2.1 Pointer position:

The pointer position is displayed in the tool bar (to the right) and it can be "fixed" by pressing the "**f**" key. This allows you to copy and paste it elsewhere.

Zone 1 (upper zone, display is the entire sound): t in seconds.

Zone 2 (window into the sound): t in seconds, and v as linear amplitude.

Zone 3 :

- On the sonogram: t in seconds and f in Hz. If the sonogram is displayed, v in dB.
- On the spectrogram: v in dB and f in Hz.

Zone 4 (sequencer tracks): t in seconds.

11.2.2 Playing back a sound from any given point:

By placing the pointer (any tool except the hand tool) in zone 1 (upper zone, displaying the entire sound) it will take the form of a loudspeaker with a vertical line.

This makes it possible, by clicking, to play back the sound from any point. The movement is shown by a red line over all zones and the position is shown in the field at the top of the AudioSculpt window.

11.2.3 Resizing a track element:

By placing any tool (except the Marker tool) at one of the track element extremities, it is transformed into two small black triangles. Then, by clicking, you can stretch or shorten the element horizontally.

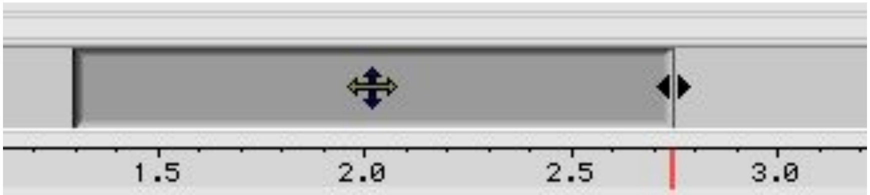

11.2.4 Moving a track element:

By placing any tool (except the Marker tool) on a track element and clicking, you can horizontally move a track element (or several track elements, if they have been selected).

11.3 The pointer (or Arrow) tool

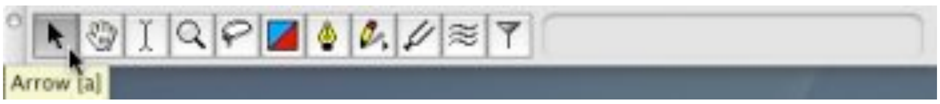

This tool has several functions, according to zones or objects:

- In zone 1 (upper zone, showing the entire sound) it changes into a small loudspeaker with

a vertical line, allowing you to play the sound from any point (please see sections 6 and 11.2.2).

- In zone 2 (window into the sound) it allows you, by clicking, to select a part of the sound.
- In zone 3 (sonogram) it allows you to move a marker (by clicking on its upper triangle). You can resize a process (the cursor changes into two small black triangles at each end of the process -- by clicking, you can change the position of the beginning or the end).
- In zone 3 (sonogram) it changes into two small black triangles when the mouse is placed on one of the rectangle's edges, or pencil stroke. This makes it possible, by clicking, to horizontally resize (the duration) or vertically resize (the frequency).
- In zone 3 (sonogram) it also makes it possible to select (by rectangle) one or several pencil strokes.
- In the zone 4 (sequencer tracks) it allows for horizontal resizing or moving of a track element (please see sections 11.2.3 and 11.2.4).

11.4 The hand tool (grab tool)

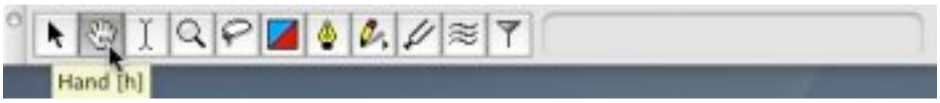

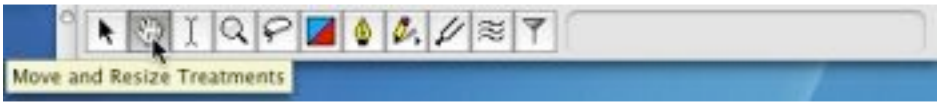

The hand tool allows you:

- to move the zooming red rectangle in zone 1;
- to grab the sonogram image and move it in any direction in its window (the " window into the sound" display. The spectrogram and track displays will follow accordingly). If you have zoomed, you can carry out this operation horizontally by grabbing the sound image in the "window into the sound";
- to grab and move any surface (form) that has been placed on the sonogram (rectangle, free and pencil stroke). To do this, select one or more of them on the sequencer tracks. If none of the surfaces (forms) have been selected, it is the sonogram image that is moved.

In zone 4, (sequencer tracks) this tool also allows you to resize or to move track elements (please see sections 11.2.3 and 11.2.4).

11.5 The cursor selection tool

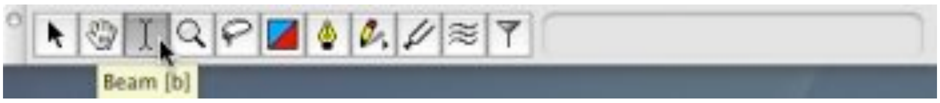

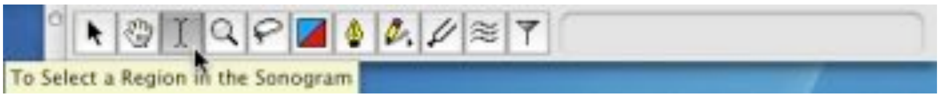

The selection cursor tool allows you to select part of the sound (simply click) in zones 2 (window into the sound) and 3 (sonogram).

Double clicking in zone 2 will select the entire sound.

This tool also allows you to: in zone 1, play the sound from any desired point (see sections 6 and 11.2.2); in zone 4 (sequencer tracks) to resize or to move the track elements (see sections 11.2.3 and 11.2.4).

11.6 The magnifying tool

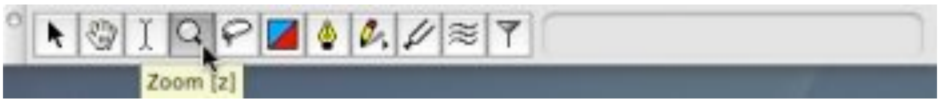

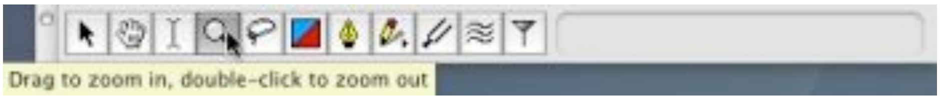

In zones 2 and 3, the magnifying tool defines the rectangle which you want to zoom: amplitude-duration (zone 2); frequency-duration (zone 3, sonogram); frequency (zone 3, spectrogram).

You can cancel a zoom by double clicking on the corresponding horizontal or vertical ruler (see section 10.2).

This tool also allows: in zone 1, play back all of the sound from any given point (see sections 6 and 11.2.2); in zone 4 (sequencer tracks) resizing or moving track elements (see sections 11.2.3 and 11.2.4)

11.7 The lasso tool

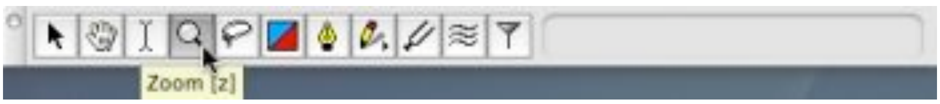

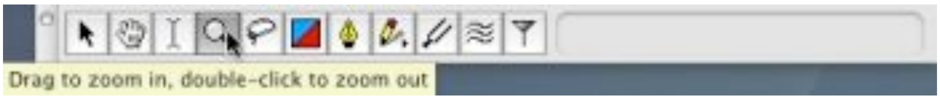

The lasso tool allows you to select several surfaces on the sonogram (you do not have to encircle them completely).

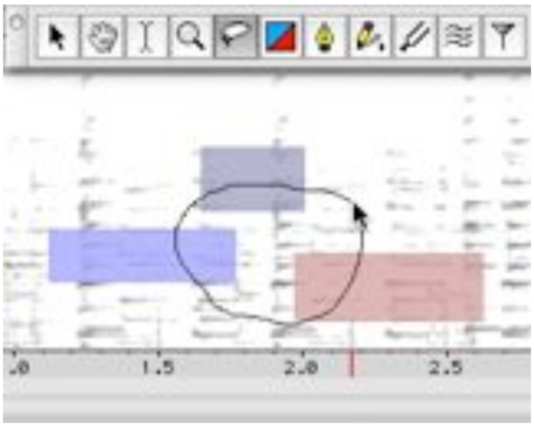

In zone 2 (window into the sound) it also allows you to select part of the sound.

This tool further allows: in zone 1, play back of the sound from any given point (see sections 6 et 11.2.2); in zone 4 (sequencer tracks) resizing or moving track elements (see sections 11.2.3 and 11.2.4).

11.8 The rectangle (rectangular surface) tool

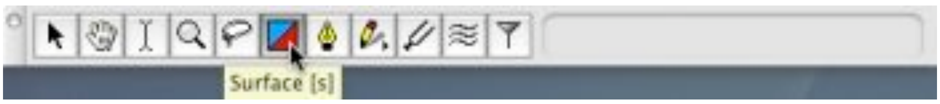

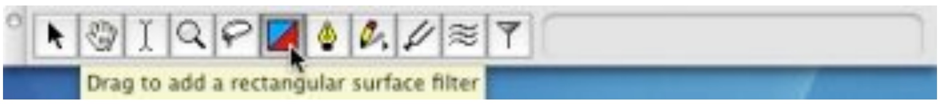

The rectangular surface tool allows you to draw a rectangular shape filtering surface on the sonogram (zone 3).

On zone 2 (window into the sound) it allows you to select part of the sound.

This tool further allows you to: in zone 1, play back of the sound from any desired point (see

sections 6 and 11.2.2); in zone 4 (sequencer tracks), it allows you resize or to move track elements (see sections 11.2.3 and 11.2.4).

11.9 The freehand surface tool

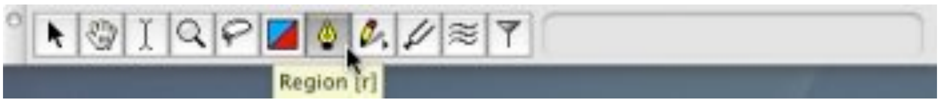

The surface tool allows you to draw a filtering surface on the sonogram (zone 3): either point by point (polygon surface) or in freehand style.

Point by point surface (polygon surface): Place the pointer at the desired point, then click to define the first point, then move the pointer (drag a segment), click to place the following point and so on. To automatically close the surface, double click.

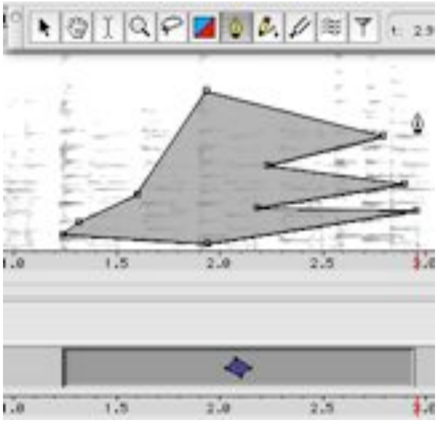

Freehand surface: first place the pointer at the desired point, then holding down the mouse button, and draw the desired outline. Close the surface automatically by releasing the mouse button.

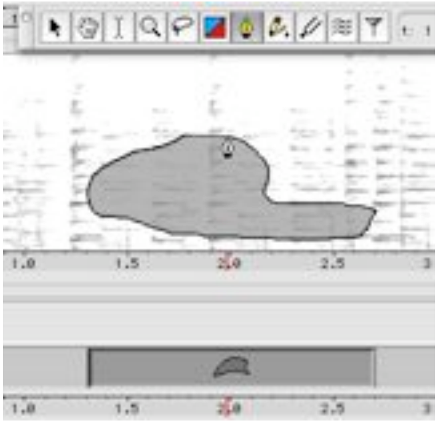

In zone 2 (window into the sound) it allows you select part of the sound.

This tool also makes it possible: in zone 1, to play back the sound from any desired point (see sections 6 and 11.2.2); in zone 4 (sequencer tracks) it allows resizing or moving track elements (see sections 11.2.3 and 11.2.4).

11.10 The pencil tool

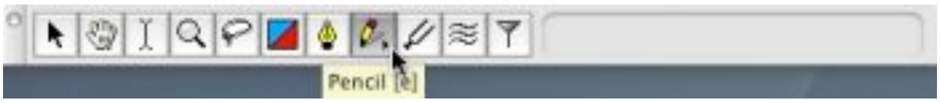

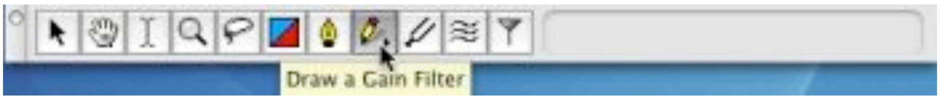

The pencil tool allows you to draw freehand over the sonogram so as to define a filtering region. You can set the thickness of the line by clicking on the small black triangle to the right of the tool in the tool box. User values can be set in the preferences section. The default thickness is 6 pixels, and the default gain is -50 dB (attenuation).

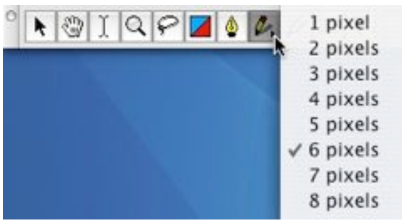

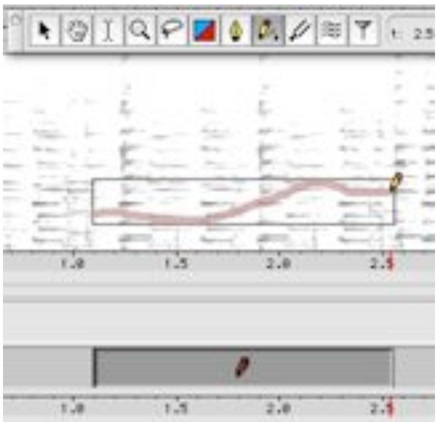

In zone 2 (window into the sound) the pencil tool allows you to select a part of the sound.

This tool also allows: in zone 1, to play back the sound from any desired point (see sections 6 and 11.2.2); in zone 4 (sequencer tracks) it allows resizing or moving track elements (see sections 11.2.3 and 11.2.4).

11.11 The tuning fork (tuning) tool

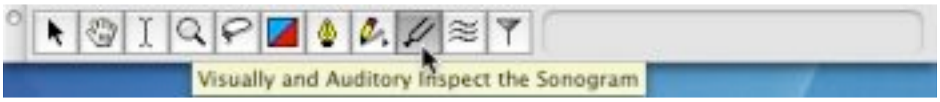

The tuning fork tool is only fully functional if the sonogram of the sound is displayed. The active part of the tuning fork is the small red ball at its base.

Amplitude, frequency and time position are shown to the right of the tool box, while the corresponding spectrogram is displayed in blue. Click to play a pure sinusoidal sound using the frequency and the amplitude that correspond to the point of the sonogram that was clicked.

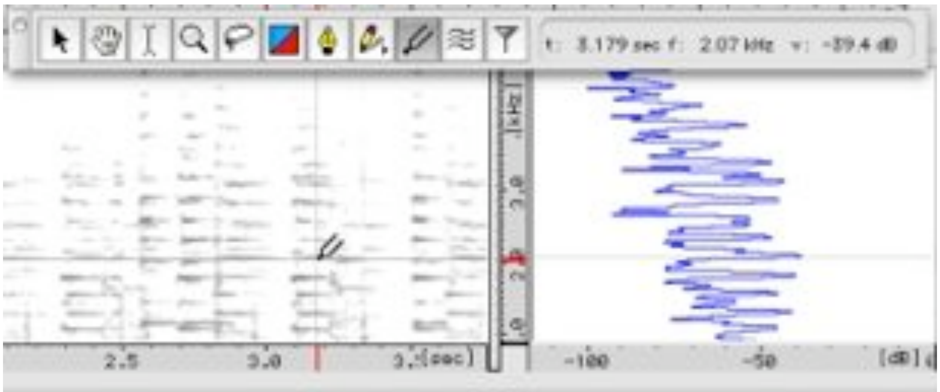

In zone 2 (window into the sound) the tuning fork tool allows you to select part of the sound.

It also allows: in zone 1, play back of the sound from any desired point (see sections 6 and 11.2.2); in zone 4 (sequencer tracks) it allows resizing or moving track elements (see sections 11.2.3 and 11.2.4) .

11.12 The harmonics tool

The harmonic tool allows you to visually set harmonics on the sonogram (this is only useful when the sonogram is displayed).

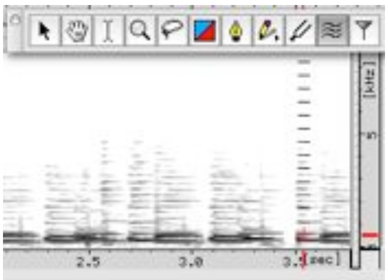

In zone 2 (window into the sound) this tool allows you to select part of the sound.

It also allows: in zone 1, play back of the sound from any desired point (see sections 6

and 11.2.2); in zone 4 (sequencer tracks) it allows resizing or moving track elements (see sections 11.2.3 and 11.2.4).

11.13 The marker tool

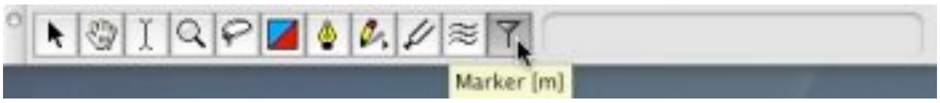

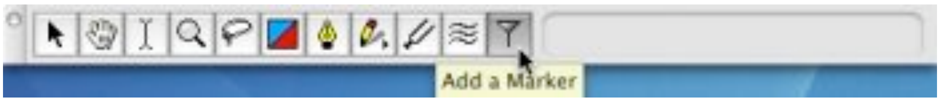

The marker tool allows you to place a marker ("Hand Added Markers") by clicking at the desired place in zone 3 (sonogram zone) as well as in zone 2 (window into the sound) or alternatively on the sequencer tracks. During play back, you can set markers on-the-fly with the keyboard shortcut '**t**' (see section 31.1).

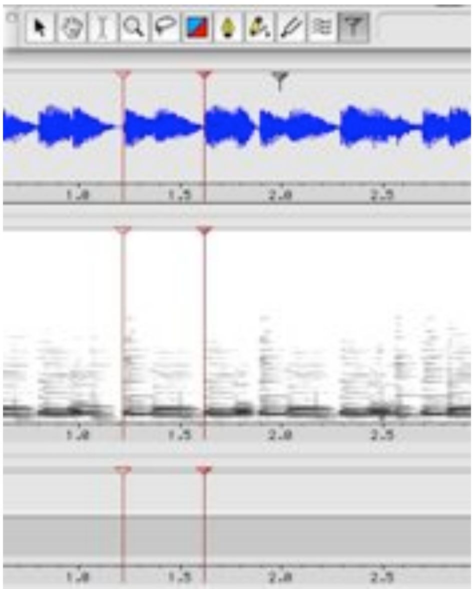

In zone 1, it allows for play back of the sound from any desired point (see sections 6 and 11.2.2).

12 The "Inspector" tool box

This floating palette only opens if the "Show Inspector" item has been checked in the "Window" menu.

The "**tab**" key toggles this palette between show/hide, as well as the three others ("Tools", "Sonogram Display", and "Grid Settings") provided they have been checked. The " **Apple +I**" ("**command +I**") keyboard shortcut toggles the "Inspector" palette between show/hide. This palette gives precise information concerning the position of the "processing" objects values, tuning and markers.

It displays information concerning a single element selected in a track. If no selection has been made, it displays only "AudioSculpt".

If several processes (or markers) have been selected, it gives information about the one that was last selected. If you require information about another process (or marker) among those that have been selected (and one that you might wish to modify) without losing the selection, just click on the desired track element (or marker) while holding down the **Shift** key.

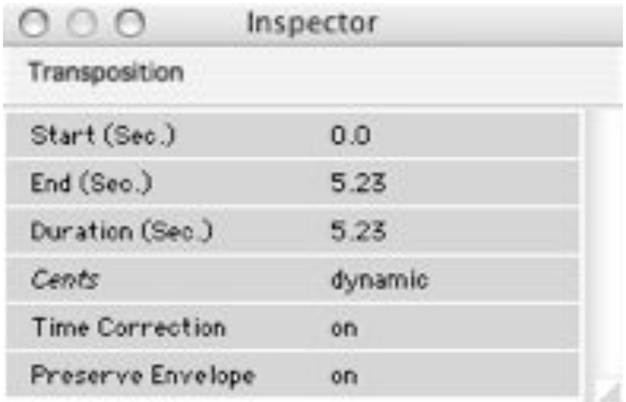

This palette is very useful when you want to control " treatment" object parameters or marker positions directly.

In order to change a parameter value, just click in its field to select it. Then type in the new value and validate with "return" or "into" (or simply by clicking elsewhere).

This makes it very simple to place and size and object with precision, by providing two parameters out of three: start time, end time, or duration.

Note: Naturally care must be taken not to type in incoherent values such as a start time that is greater than an end time. Such a modification would be refused.

This palette gives access to all object parameters. However some of them are of course not changeable: these are displayed in italics. For example, in Bpf treatments, the variable is marked "dynamic" and cannot be changed.

Some fields can only be "reversed" between on and off by means of a click ("Time Correction" in the case of a transposition, for example).

Note: The "Inspector" palette only gives information about a single object, and allows changes to that object only.

13 The "Sonogram Display" palette [toolbox]

This floating palette only opens if the "Sonogram Display" item has been checked in the "Window" menu.

The **tab** key toggles this palette between hide/show, as well as the three others (Tools, Inspector and Grid Settings) on condition they have been checked.

The "**Apple +J**" ("**command +J**") keyboard shortcut toggles the "Sonogram Display" palette between show/hide.

It toggles the display between show/hide in zone 2 (the sonogram zone). What it displays depends on the analyses that have been carried out as well as the presence of manually added markers and processes.

If no analysis has been carried out, if no process has been defined, if no marker has been created manually, then it displays nothing.

Each box (and if applicable their associated horizontal cursors) will only be displayed if the corresponding element has been calculated or defined.

When a sonogram is displayed, there are 2 horizontal linear sliders for setting its grayscale display (see section 16.2):

The default settings are: - 10 dB for black ("Black Threshold") and - 60 dB for white ("White Threshold").

Each marker type generated during analysis possesses a slider for setting the display threshold (see section 17).

14 Cursor position

If no selection has been made the default cursor position will be at the start of the sound. If a selection has been made, the default position is at the start of the selection. The cursor may be in any position if play back is interrupted. At the top left of the main window, the first field displays cursor position (during play back or when stopped).

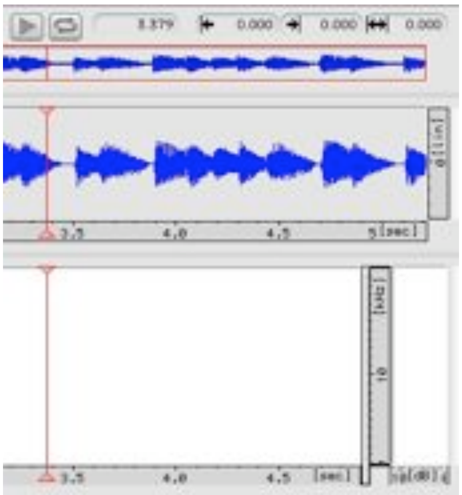

However, it is easy to choose cursor position.

Just grab the upper triangle of the cursor using any tool (except the magnification tool, the tuning fork to or the marker tool, see section 11). Click and move it to the desired point.

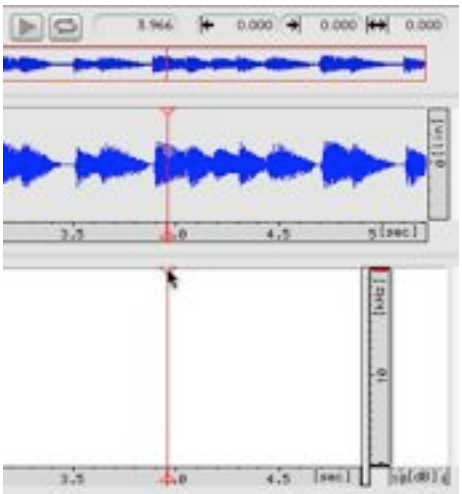

Alternatively, you may type the time position value into the first changeable field, in the upper right hand part of the AudioSculpt window. Validate with "into" or "return" or simply by clicking elsewhere.

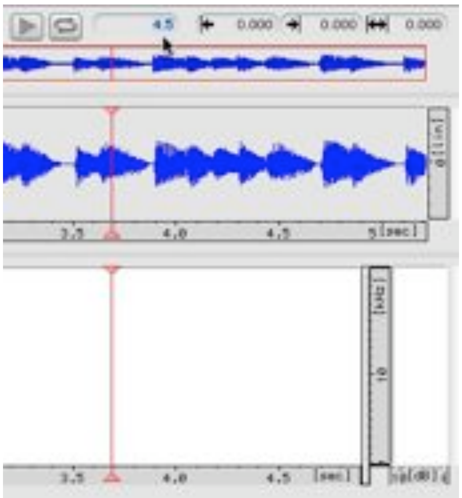

Using the arrow tool, just click at the desired point in the "window into the sound" zone (zone 2). Any previously made selection will not be affected by the change in cursor position.

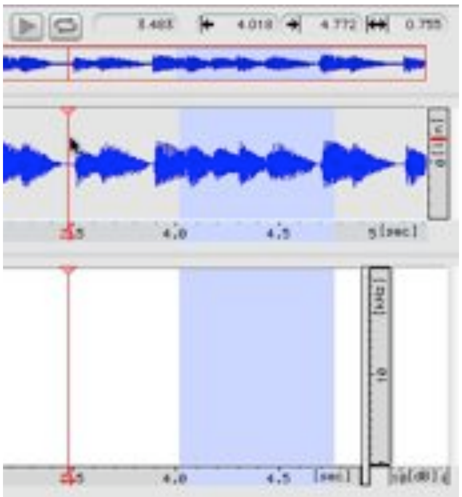

Using the selection cursor (see section 11.5), just click at the desired point in the sonogram the zone (zone 3). However in this case you lose the selection (the selected zone is de-selected).

15 Selecting part of the sound

Selecting is usually carried out using a tool (see section 11). The dedicated tool for this is the selection cursor tool which can be used in zone is 2 (window into the sound) and 3 (sonogram zone). To select the entire sound in zone 2, just double click.

The following tools: Arrow, lasso, rectangular surface, free surface, tuning fork, as well as the harmonic tool, may all be used in the same way, but only in the window into the sound (zone 2). Selection values (as well as cursor position) are displayed in the upper right hand part of the AudioSculpt window.

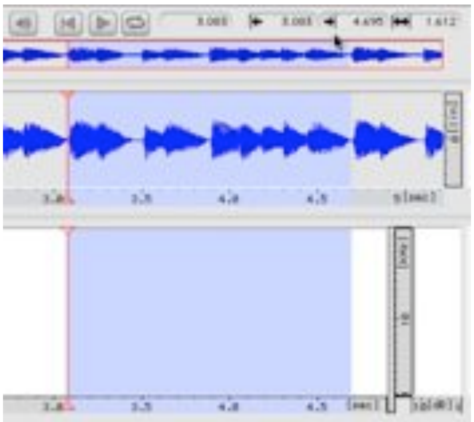

From left to right, the 4 fields display:

- cursor position (in play back or stopped);
- selection start and end (if there is a selection);
- duration of the selection (if there is a selection)

All these fields may be edited. Therefore the selection may be defined by typing in 2 values out of 3 (start, end, duration). To edit a parameter value, simply click to select it, then type in the new value and validate via **return** or **enter**, or simply clicking elsewhere.

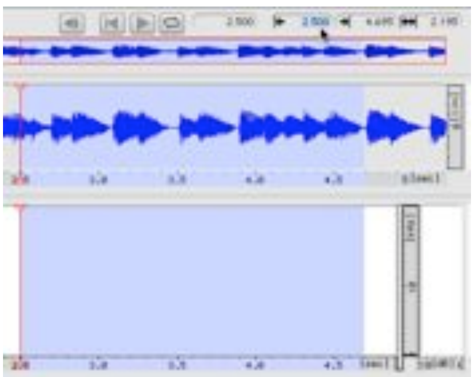

To deselect all or part of the sound, simply:

- make a new selection;
- place the cursor tool at any point in the sonogram zone and click.

16 Analysis

Analysis can only be carried out on open sound files.

16.1 General analysis parameters

These parameters may be set in the panel area that is opened when certain analyses are carried out ("Analysis Settings", "FFT Settings"). They deal with Sonogram Analysis, Fundamental Frequency Analysis (F0), Chord Sequence Analysis, and analysis generated markers ("Generate Markers. . . ").

- **"Window Size"** : window size determines the number of sound samples in each analysis. Window size is the main parameter in analysis: it determines frequency and time resolution.

  The "**Window Size**" and "**Fundamental Frequency**" fields are interdependent.

- **"Fundamental Frequency" in Hz** : allows direct adjustment of the frequency resolution. Window size, step resolution and FFT are adjusted accordingly. Window size is programmed to be 5 times as long as the fundamental frequency period shown (frequency resolution).
- **"Window Step"** : determines the time interval between two successive analyses. This is measured in samples. In manual mode, set the window step manually (for advanced users: flag -I).

  In automatic mode: window step will be adjusted automatically according to the required operation, and to obtain the best result. In analysis, the window step is equal to 1/8th of the window size.

- **"FFT Size"** : determines the number of analysis points; this is necessarily equal to or greater than window size.
- **"Analysis Window"** : specifies the window type used in the analysis. Three window types are available: Blackman, Hanning and Hamming.

  Note: by default, the "Verbose Output" box is unchecked. If you check this option (for advanced users: flag -v), all the information provided by SuperVP, except for the command line, will be displayed in the SuperVP console. The information will only be visible if the console is folded out (see section 24.1).

16.2 "Sonogram analysis" : the sonogram

For sounds that are not mono this analysis is carried out on one of the channels, chosen by the user, or upon all of the channels simultaneously.

Three analysis types are available:

- Fourier Transform analysis: FFT;
- Linear Prediction analysis: LPC;

- Discrete Cepstrum analysis.

Choose "Sonogram Analysis. . . " in the "Analysis" menu:

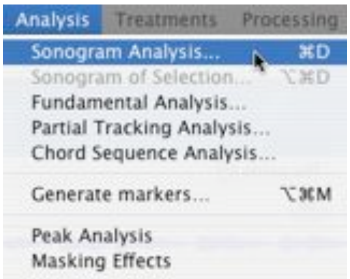

If part of the sound is selected, you can choose "Sonogram of Selection. . . ": only that part of the sound will be analyzed. The corresponding sonogram will be displayed.

A dialog box will open to allow you to choose your parameters:

In addition to the above described parameters, it is necessary to determine:

- **"Analysis Type"** : this allows you to choose between FFT, LPC and Discrete Cepstrum.
- **"LPC Order"** : this defines the number of LPC totals or the number of Cepstrum coefficients for Discrete Cepstrum.
- **"Channel to analyse"** : this allows you to choose either one or all of the channels (for sounds that are not mono).

Then, just click on "Do Analysis".

At the bottom of the AudioSculpt window, in the console area (zone 5), the appropriate information and progress bar will be displayed.

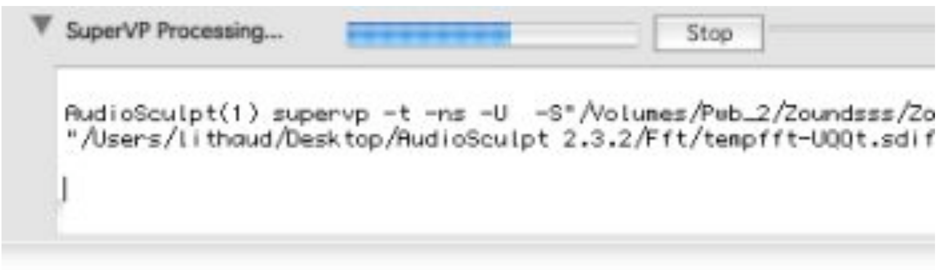

When this is finished, the sonogram will be displayed.

The "Stop" button at the right of the progress bar allows you to suspend or interrupt the process: a partial sonogram will be displayed.

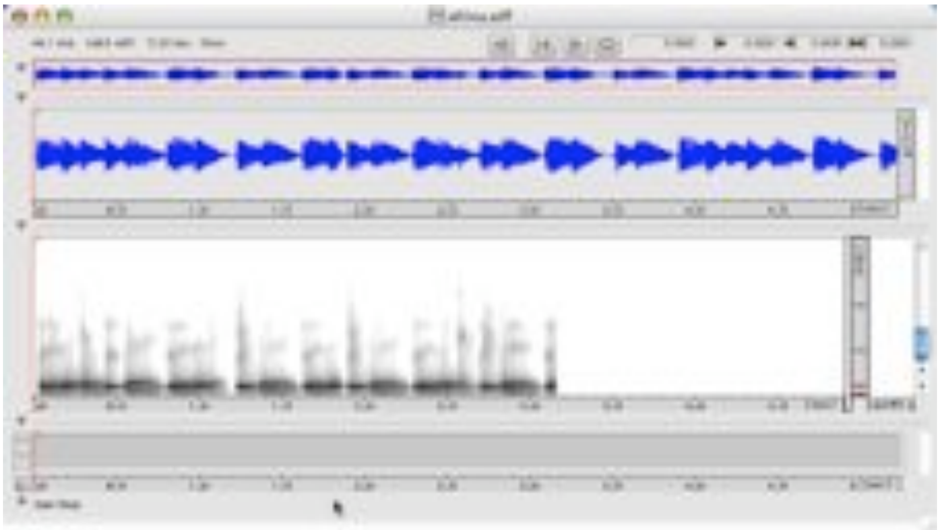

The default frequency band displayed lies between 0 and 7000 Hz, where the most interesting elements reside. Double click on the vertical scale to display the entire sonogram (between 0 and 22 050 Hz, for a sound sampled at 44 100 Hz).

You can change sonogram display via the "Sonogram Display" floating palette.

> **Note:** This will only open if you have checked the " Show Sonogram Display" in the "Window" menu (see sections 13 and 31.6).

The **tab** key toggles this palette between hide/show, as well as the three others (" Tools", "Inspector", "Grid Settings") on condition they have been checked.

The keyboard shortcut "**Apple +J**" ("**command +J**") toggles the "Sonogram Display" palette between hide/show.

You can deactivate the sonogram display by unchecking the "Sonogram" box. It is checked by default.

Two horizontal sliders allowing you to adjust the grayscale display.

The default settings are:

- "Black Threshold" : -10 dB, for black
- "White Threshold" : -60 dB, for white.

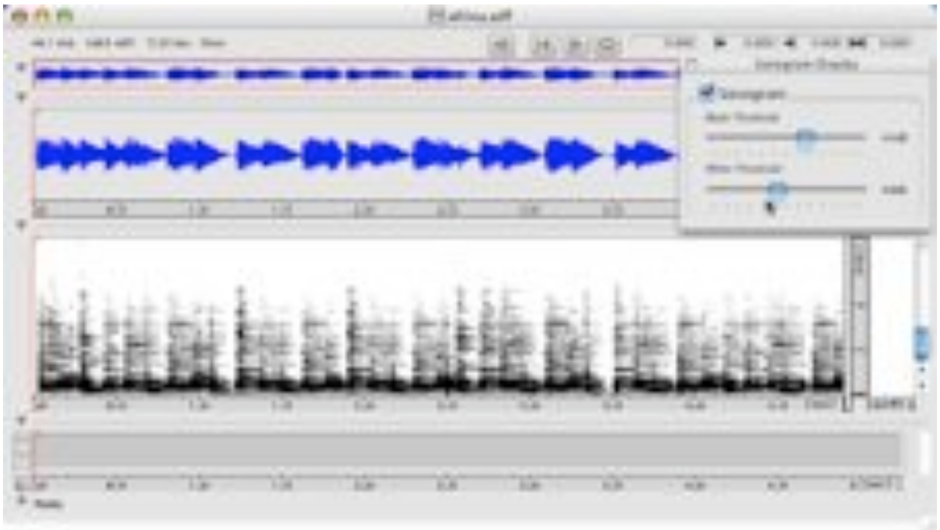

In the "Window" menu, the "Optimize Sonogram" item allows you to run full-screen mode for the sonogram and the "window into the sound" (zones 2 and 3), and return to the initial size. The keyboard shortcut is "**Apple +U**" ("**command +U**").

Sound sonograms may be of any duration. In addition to processing, represented by a progress bar, display may also take time: in this case the display percentage (progress indication) is shown on the sonogram. The same applies to multichannel sounds (several sonograms are displayed).

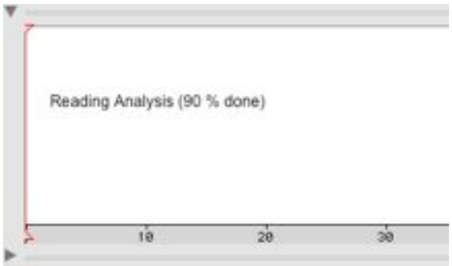

After zooming into the sonogram of a long sound, you can navigate around the sound. However after a certain time, AudioSculpt will need to load a new detailed image. Here the display percentage (progress indication) will also be shown.

The "Export Image" function in the "File" menu allows you to export and image of the sonogram and/or the wave form under PNG, TIFF, or JPEG format.

Important: The image generated may be very large.

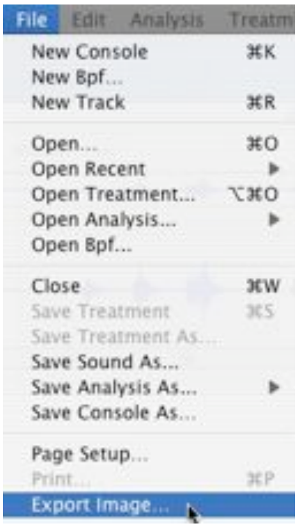

Saving the SDIF format file containing the sonographic image is done via " Safe Analysis As. . . " in the "File" menu. Simply choose " Save Sonogram As. . . " when the submenu comes up.

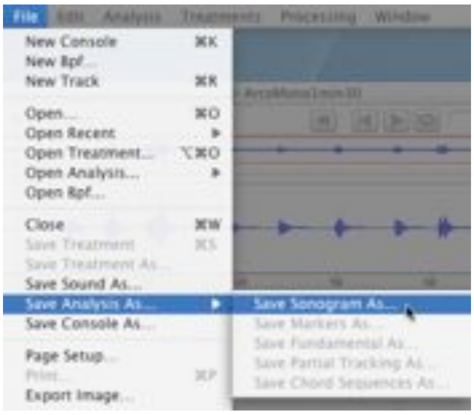

A standard dialog window will ask you to provide a name (by default "name_of_sound.fft.sdif"), as well as the location where you want to save the file. The default location is the "Fft" folder, in the application folder, but saving can be done at any other location.

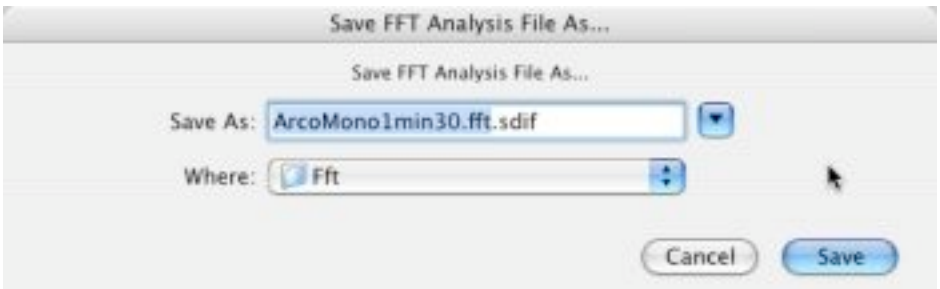

Thus you can save several sonograms of the same sound, carried out with different parameters.

> **Note:** no matter what the type of the sonogram carried out (FFT, LPC or Discrete Cepstrum) the default name is always "name_of_sound.fft.sdif". It is advisable to type in a more appropriate name.

The "Open Sonogram. . . " in the "Open Analysis. . . " submenu in the "File" menu will open a standard dialog box allowing you to choose from among previously saved files.

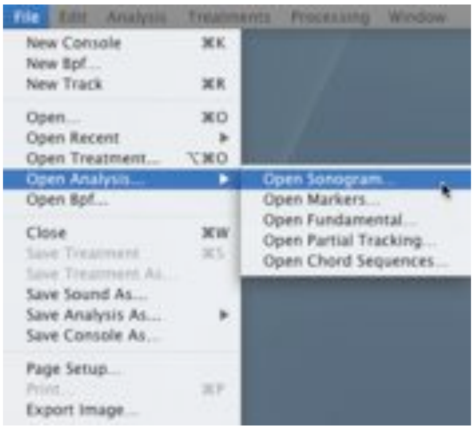

> **Note:** if a sound file is open, and you request that an analysis file be opened concerning a different sound, AudioSculpt will ask for confirmation. By default, it will always start by searching and opening the sound file that corresponds to the analysis file.

> **Note:** if no sound file is open, and you request that an analysis file be opened, AudioSculpt will open the corresponding sound file.

16.3 "Fundamental Analysis": fundamental frequency (F0) analysis

This analysis looks for the fundamental frequency by analyzing the spectral peaks. An estimation of the fundamental frequency is carried out in each analysis window.

Analysis is carried out upon part or all of the sound. For sounds that are not mono, analysis

is carried out upon a channel chosen by the user, or upon all of the channels simultaneously.

The analysis will generate a SDIF file containing the estimation of the sound's fundamental frequency. This F0 estimation is displayed as a Bpf in the sonogram window.

If the sonogram is currently displayed, the estimation will be superimposed upon it. This Bpf can be edited and changed in the Bpf editor (to see section 19.4.1).

The "Fundamental Analysis. . . " item in the "Analysis" menu opens the settings panel for "Fundamental Analysis Parameters". Here you set the usual analysis parameters of (see section 16.1). You also set the following:

- **"Fundamental Minimal Frequency" in Hz** : this is the minimum fundamental frequency threshold below which a fundamental frequency will not be looked for. The default is 50 Hz.
- **"Fundamental Maximal Frequency" in Hz** : this is the maximum fundamental frequency threshold above which a fundamental frequency will not be looked for. The default is 1000 Hz.
- **"Maximal Frequency in Spectrum" in Hz** : the analysis will not take into account spectral peaks about this threshold. The default is 4000 Hz.
- **"Noise Threshold" in dB** : this specifies a noise level: if the amplitude difference between a given peak and the highest peak is greater than this value, the peak in question will not be taken into account.
- **"Smooth Order"** : frequency smoothing filter order. The default is 3.
- in addition, the selection may be kept within certain limits, or another selection may be defined. Just check the "Restrict to Selection" box, and if applicable, type in new values in the appropriate fields.
- **"Channel to analyse"** : this allows you to choose either a single channel, or all channels (for sounds that are not mono).

The "Do Analysis" button will start the processing. Progress will be shown by a progress bar displayed at the top of zone 5 (integrated console). When processing is done, the Bpf will be displayed on the sonogram.

Note: processing may be stopped at any time by pressing the "Stop" button located at the right off the progress bar. In this case, the partial result will be displayed.

You can show/hide the display in the "Sonogram Display" (see section 13).

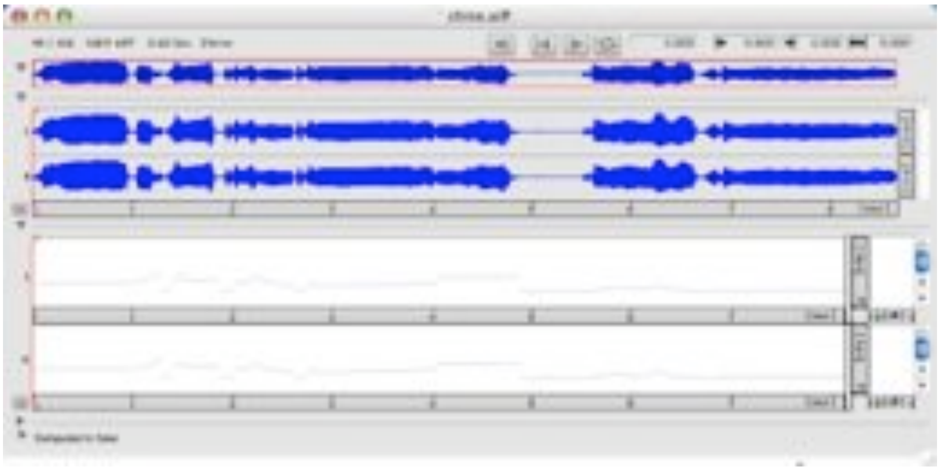

To erase the Bpf, click on the curve and press the "**Delete**" button "<-" or "**Backspace**" key).

To edit it, double click on the curve. The Bpf editor will open, and you can carry out editing using the pencil or the eraser tool in the "Bpf Tools. . . ;" menu (please consult the Diphone Studio documentation).

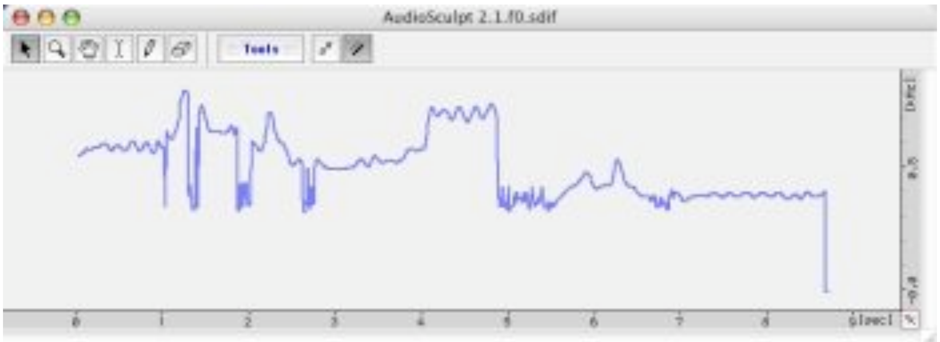

When you close the modified Bpf, at dialog box allows you to save the modifications that you carried out to the SDIF file. However this will replace the previous file which has the same name.

The modifications will be shown on the sonogram.

Several F0 files may be saved in the course of the session (the "Save Analysis As. . . " in the "File" menu, "Save Fundamental As. . . " submenu).

However, be careful, as "Save Fundamental As. . . " only saves the last file (the same applies to the save panel that comes up when the sound file is closed; see section 29).

Analysis parameters may be changed for part of the sound. The other parts of the sound will continue to use the old parameters.

> **Note:** you cannot save F0 type Bpf's in the same way that you save the other Bpf's. If you wish to recover it in order to reuse it as a Bpf, you will need to copy it and paste it in a new open Bpf via "New Bpf" under the "File" menu.

You can save the corresponding SDIF file in the usual way (see section 29). To save this file, use "Save Fundamental As...", submenu of "Save Analysis..." in the "File" menu. The default name is "my_sound.f0.sdif". The file will be exploited from any location. The default location is the "Fundamental" folder, located in the application folder, but you may save the file at any desired location.

To open previously saved files, use "Open Fundamental..." in the "Opened Analysis..." submenu in the "File". A standard dialog box will open.

> **Note:** if a sound file is open, and you request that an analysis file be opened concerning a different sound, AudioSculpt will ask for confirmation. By default, it will always start by searching and opening the sound file that corresponds to the analysis file.

> **Note:** if no sound file is open, and you request that an analysis file be opened, AudioSculpt will open the corresponding sound file.

16.4 Partial Tracking Analysis

Partials estimation and tracking is carried out in each analysis window.

This analysis to be carried out on part or all of the sound. the selection may be kept within certain limits, or another selection may be defined. Just check the "Restrict to Selection" box, and if applicable, type in new values in the appropriate fields. The analysis generates a SDIF file that contains the sound file partials tracking estimate. This estimate the displayed as multi-Bpf's in the sonogram window.

If the sonogram is already displayed, the analysis will be superimposed over it. The multi-Bpf may be edited and changed in the Bpf editor (see section 19.4.1).

> **Important:** for stereo and multichannel sounds, partials tracking analysis can only be carried out on one channel at a time.

- **"Channel to analyse"** : this item allows you to choose the channel you wish to analyze (for sounds that are not mono).

If you stop the processing, analysis is canceled and SuperVP (in fact pm2) will return an error message.

16.4.1 Inharmonic analysis

In the "Analysis" menu, use the "Partial Tracking Analysis...;" to open the "Partial Tracking Parameters" control panel. Here, you set the usual FFT analysis parameters (see section 16.1) as well as the general in harmonic analysis parameters (displayed by default).

- **"Maximum Number of Partials"**
- **"Minimum Amplitude Level" en dB**

- **"Smoothing Envelope Attack"** in seconde : duration of the attack slope (at the present time, this value must be equal to zero for the analysis).
- **"Smoothing Envelope Release"** in seconde : duration of the release slope (at the present time, this value must be equal to zero for the analysis).

"Peak Connection" : sets the connection between the peaks contained in two successive analysis windows. The peaks to be connected are selected according to the allowed frequency variation and amplitude variation between them.

Allowed frequency variation comprises 2 components: a component expressed in relative frequency and a component expressed in absolute frequency.

- **"Relative Frequency Deviation"** expressed in hundredths of a semitone ("cents"): variance of the relative frequency distance. The default value is 20 semitone hundredths ("cents").
- **"Constant Frequency Deviation"** in Hz : absolute frequency distance variance. The default value is zero.
- **"Relative Amplitude Deviation"** : in percent: relative distance variance between amplitudes. The default value is 50%.
- **"Source Partial Neighbors"** (between 1 and 10): the number of neighboring peaks considered as possible source candidates for local connection optimizing. The default value is 1.
- **"Target Partial Neighbors"** (between 1 and 10): the number of neighboring peaks considered as possible target candidates. The source value must be less than or equal to the target value. The default value is 3.

"Partial Connection" : sets the connection between partial fragments.

- **"Time Gap to Connect Over"** in seconds; maximum Time distance between two fragments to be connected. The default value is 0 .017 seconds.
- **"Frequency Gap to Connect Over"** in semitone hundredths ("cents"): relative maximum frequency distance between two fragments to be connected. The default value is 20 semitone hundredths
- **"Minimum Partial Length"** in seconds: minimum duration of a partial; shorter fragments are eliminated. The default value is 0.009 seconds.

16.4.2 Harmonic analysis

The default setting is inharmonic analysis. To carry out "Harmonic Analysis" the option must be checked (parameters will then be available).

First it is necessary to carry out a fundamental frequency analysis (f0), to save it (it is advisable to use an appropriate name, for sample "my_sound.f0.sdif"). Then, select it using the "Select" button.

- The "Select" button opens a standard dialog box allowing you to choose a "my_sound.f0.sdif" type file.
- "**Bandwidth Partial Seeve**" : an f0 coefficient multiplier that defines the bandwidths in which harmonics must be looked for. This is a value between 0 and 1; for a value 1, bandwidth will be equal to f0.

> **Important:** the parameters defined in " Peak Connection" and "Partial Connection" are not activated.

The "Do Analysis" starts the processing, which you can follow thanks to the progress bar at the top of the zone 5 (integrated console). The analysis generates a SDIF file containing the sound file partials tracking estimate. These partials are displayed (in color) in the sonogram window. You can hide the display in the " Sonogram Display" palette. (See section 13).

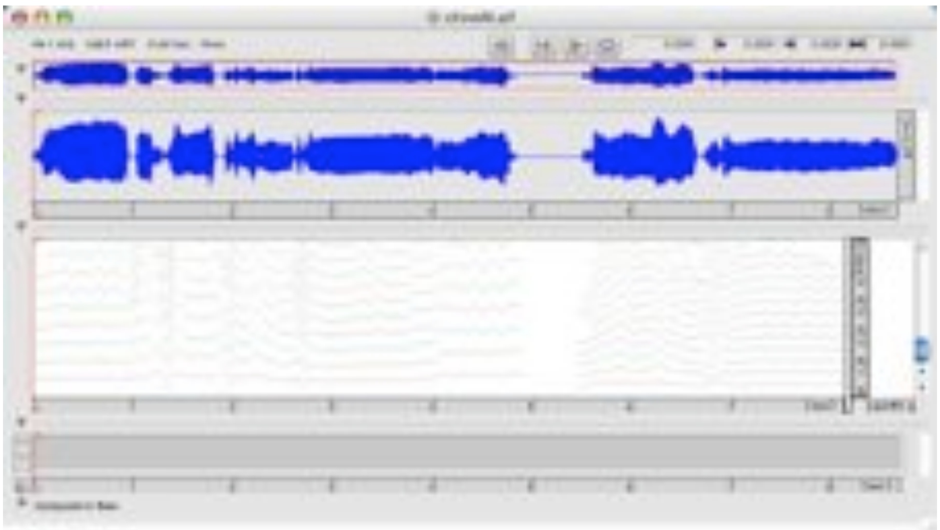

The file lasts obtained can be saved via "Save Partial Tracking As. . . " in the "File" menu (the default name is "my_sound.hrm.sdif" for harmonic analysis, and "my_sound.trc.sdif" for in harmonic analysis). The file may be used elsewhere.

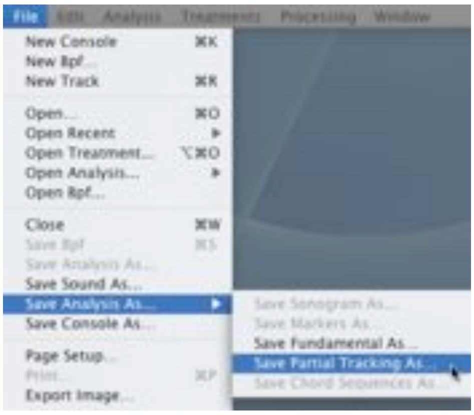

In order to erase the analysis results, click on one of the curves and carry out "**Delete**" ("<**-**" or "**Backspace**" key).

In order to edit them, double click on one of the curves. The Bpf editor will open and will allow changes via the pencil, the eraser, and "Bpf Tools" (please consult the Diphone Studio documentation).

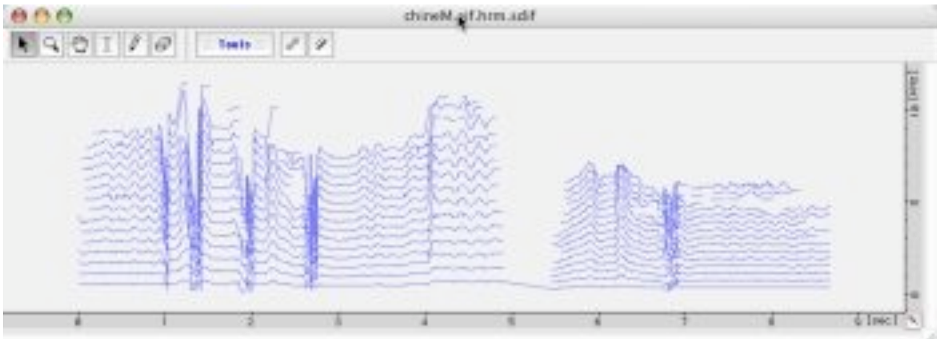

When you close the modified Bpf, a dialog box will ask if you want to save changes to the SDIF file. Note that it will replace the previous file of the same name. Use the non item "Save Analysis As. . . " in the "File" menu to change the name. The sonogram will be displayed with the modifications.

Several files may be saved in the course of the session (the "Save Partial Tracking As. . . " in the "File" menu, "Save Fundamental As. . . " submenu). However, be careful, as "Save Partial Tracking As. . . " only saves the last file (the same applies to the save panel that comes up when the sound file is closed; see the "File saving" section). The "Open Partial Tracking. . . " in the "Open Analysis. . . " submenu in the "File" menu will open a standard dialog box allowing you to select a previously saved file.

> **Note:** if a sound file is open, and you request that an analysis file be opened concerning a different sound, AudioSculpt will ask for confirmation. By default, it will always start by searching and opening the sound file that corresponds to the analysis file.

> **Note:** if no sound file is open, and you request that an analysis file be opened, AudioSculpt will open the corresponding sound file.

16.5 Chord Sequence Analysis

In a region, a "chord" is a fixed frequency and amplitude partials group represented by straight horizontal segments, the beginning and end of which are defined by region markers. This analysis generates a SDIF file containing the sound file partials tracking estimate. The estimate is displayed as multi-Bpf's in the sonogram window. If the sonogram is already displayed, the analysis results are superimposed over it.

Thus, chord sequence analysis is carried out on one or several regions defined by markers. In order to do this, the "And Chord Seq Markers" (keyboard shortcut "**Q**") in the "Edit" menu lays down a pair of markers, a start marker and an end marker at each end of the selection. This makes it possible to define several non-neighboring regions.

You may also use any type of marker (manually added or analysis generated; see section 17). The minimum required number is 2. The first marker will be the start marker, and the last marker will be the end marker. Intervening markers will be considered as both end markers for the preceding region and start markers for the following region.

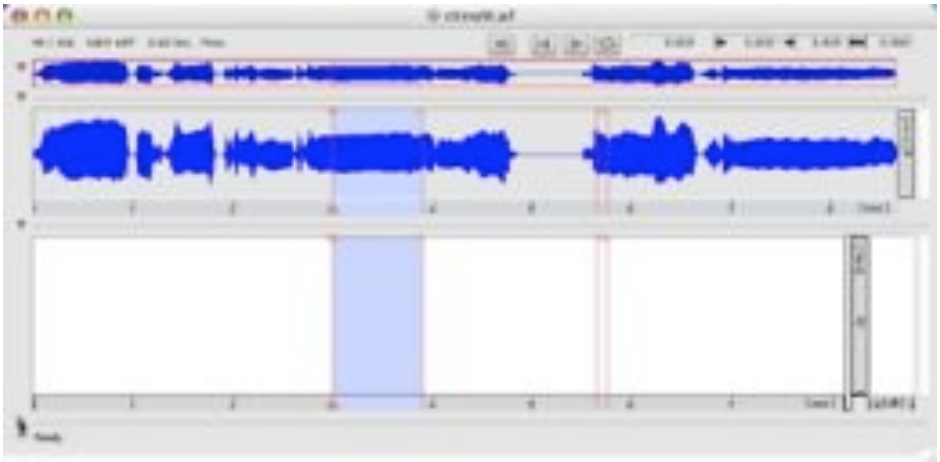

The "Chord Sequence Analysis. . . " in the "Analysis" opens the parameter settings panel.

Note: "Restrict to Selection" makes it possible to limit processing to part of the sound, on condition that it contains at least 2 markers. If you have selected part of the sound, the box will be checked by default. Remember to lay down markers at either end of the selection.

There are two analysis methods:

- Averaged Spectrum
- Inharmonic Partial Averaging.

16.5.1 Average Spectrum

This analysis calculates the Spectrum average over the segment in question (between two adjoining markers) and lines up a partial with each peak in accordance with the parameters below.

The ("FFT Settings") are the usual analysis parameters

- **"Maximum Number of Partials"**: determines the maximum number of partials to be looked for.
- **"Maximum Amplitude Level (Att.)"** : the maximum amplitude attenuation value based on 0 dB. In fact this corresponds to the minimum amplitude value of peaks that can be considered as partials.

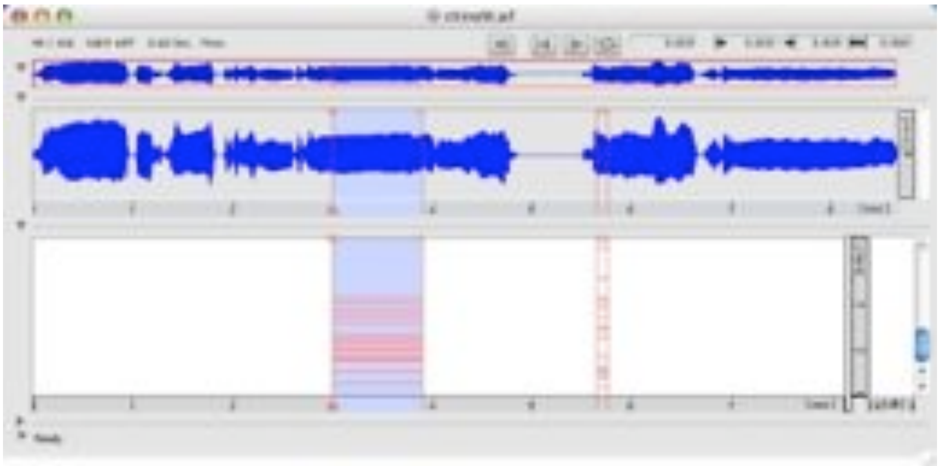

16.5.2 Inharmonic Partial Averging

This analysis calculates the partial average obtained via partial tracking analysis in inharmonic mode. A partial is considered a chord component only if it is present over the entire segment lying between two adjoining markers. In order to get satisfactory results, you will have to adapt the parameters in the "Partial Tracking" inharmonic analysis in such a way that some partials obtained stretch out over the entire segment. You must therefore carry out the "Partial Tracking Analysis in inharmonic mode in order to get visual representation of partial behavior. If the result appears promising, you then simply need to apply the same parameters to the Chord Sequence Analysis.

Note: During Chord Sequence Analysis, the "**Smoothing Envelope Attack**" and " **Smoothing Envelope Release**" are set at zero. Therefore it is important that they be also set at zero for the "**Partial Tracking Analysis**" preparatory analysis in inharmonic Mode.

Note: At the present time, it is impossible to export these parameters: they have to be noted and keyed in manually.

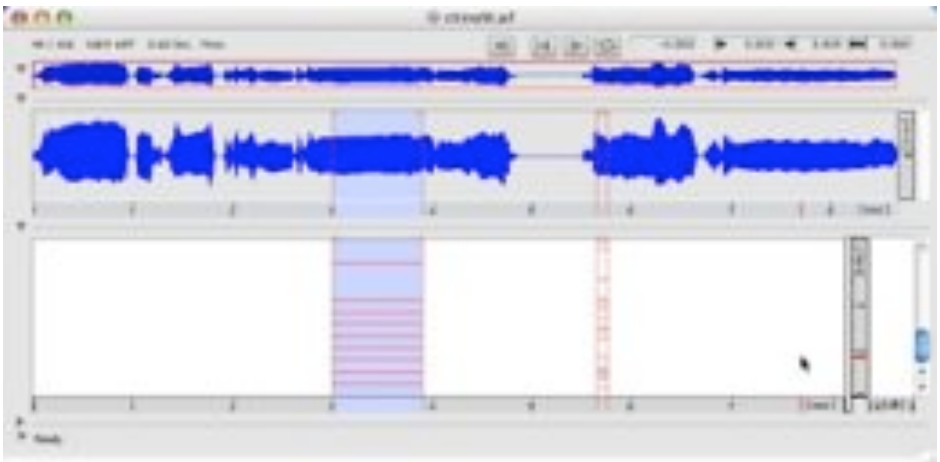

Important: for stereo and multichannel sound, chord sequence analysis can only be carried out on one channel at a time.

Several files (with different parameters) may be saved in the course of the session. However, be careful, as "Save Chord Sequences As. . . " only saves the last file (the same applies to the save panel that comes up when the sound file is closed; see the "File saving" section). The "Open Chord Sequences. . . " in the "Open Analysis. . . " submenu in the "File" menu will open a standard dialog box allowing you to select a previously saved file.

Note: If a sound file is open, and you request that an analysis file be opened concerning a different sound, AudioSculpt will ask for confirmation. By default, it will always start by searching and opening the sound file that corresponds to the analysis file.

Note: If no sound file is open, and you request that an analysis file be opened, AudioSculpt will open the corresponding sound file.

16.6 Analysis generated markers

There are three types of markers that can be generated by analysis:

- **Transient detection** markers, using the SuperVP transient detection algorithm.
- markers known as "**Spectacle Flow Markers**" (detection of rapid energy variation in the signal).

There are two options:

- "**Spectral differencing** (positive)" which detects rapid energy increase in the signal (and which may be linked to a note start);
- "**Spectral differencing** (negative)" which detects rapid energy decrease in the signal (and which may be linked to a note end).

Please refer to section 17.

The "Open Markers. . . " in the "Open Analysis. . . " submenu in the "File" menu will open a standard dialog box allowing you to choose from among previously saved files.

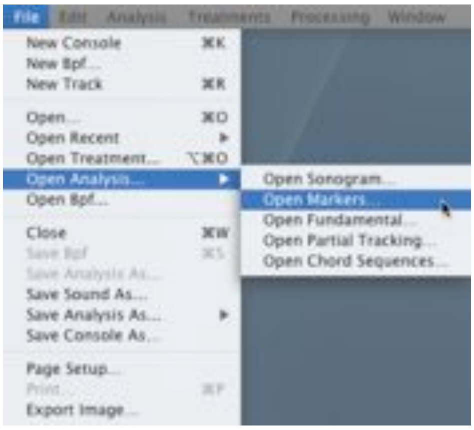

Note: If a sound file is open, and you request that an analysis file be opened concerning a different sound, AudioSculpt will ask for confirmation. By default, it will always start by searching and opening the sound file that corresponds to the analysis file.

Note: If no sound file is open, and you request that an analysis file be opened, AudioSculpt will open the corresponding sound file.

16.7 Peak Detection

This analysis looks for maximum values in adjoining short term amplitude spectra.

The analysis is applied to the entire sound. The results are not displayed in the sonogram, but are written directly into a text file (see sections 29 and 30.1).

The "Peak Analysis" item in the "Analysis" menu opens the settings panel for "Peak Detection". Here you set the usual analysis parameters, as well as:

- **"Threshold" in dB** : amplitude differences between the crest of the Peak and neighboring minimum values must be above this threshold, otherwise the peaks will not be taken into account. The default value is 25 dB. If the threshold is equal to zero, all peaks will be taken into account.
- **"Number of Peaks"** : allows you to specify the number "n" of peaks that must be taken into account in each analysis window. Only "n" peaks with the highest amplitudes will be displayed in the results. The default value for "n" is 3.

Note: By default, the "Verbose Output" box is unchecked. If you check this option (for advanced users: flag -v), all the information provided by SuperVP, except for the command line, will be displayed in the SuperVP console. The information will only be visible if the console is folded out (see section 24.1).

Validate by pressing "OK". The window will close, and a standard dialog window will open, allowing you to type in the name and location where you want to save the result. After having done this, validate with "OK". The window will close, and processing will begin.

The results are written into a text file:

- the first line shows the time of analysis: t in seconds,

- the following lines (depending upon the number of the peaks taken into consideration): frequency (in Hz), amplitude (in dB) of the peaks, arranged in order of increasing frequency. The data is laid out in columns separated by tabulations.

16.8 Masking Effects

Terhardt's algorithm reduces the amplitude spectrum to the small number of peaks that contradict to the overall perception of the sound. In fact, it models the inner ear's mask effect filtering.

The analysis is applied to the entire sound. The results are not displayed in the sonogram, but are written directly into a text file (see sections 29 and 30.1).

In the "Analysis" menu, "Masking Effects" will set the usual masking effects analysis parameters, as well as:

- **"Threshold" (threshold in dB)** : the amplitude difference between the peak located at index i and neighboring peaks located at i-3, i-2, i+2, i+3, must not exceed this threshold, otherwise the peaks will not be taken into account. The default value is -60 dB.
- **"Maximum Number of Peaks"** : allows you to specify the number "n" of peaks that must be taken into account in each analysis window. Only "n" peaks with the highest amplitudes will be displayed in the results. The default value for "n" is 3.

> **Note:** By default, the "Verbose Output" box is unchecked. If you check this option (for advanced users: flag -v), all the information provided by SuperVP, except for the command line, will be displayed in the SuperVP console. The information will only be visible if the console is folded out (see section 24.1).

The following results are stored in the file:

- **"Time and/or Number"** : all of the peaks taken into consideration.
- **"Frequency" (optional)** : Peak frequency.

- **"True Pitch"** : provides an estimate of "True Spectral Pitch" : spectral component pitch is the same as that perceived by the listener.
- **"Amplitude" (optional)** : Peak amplitude.
- **"Weight" (optional)** : This measures the spectral size of components. It is always greater than zero.
- **"Excess Sound Pressure Level"** : This value is provided for each peak and indicates its level of perception above the mask threshold. The higher it is, the more perceptible the sinusoidal value attached to the peak. A value of zero means it cannot be perceived.

Frequency, amplitude and weight are optional. The order and the choice of storage location concerning time and number of peaks is intended to facilitate the use of this kind of file in "OpenMusic".

In the text file, the first line shows the point in time at which the analysis is carried out and/or the number of peaks (according to the option selected). The following lines show the required data separated by tabulated columns.

17 Markers

Several marker types are available:

17.1 Hand Markers

The marker tool (see section 11.13) allows you to set a mark at any point in the sonogram, and also in zone 2 (sound window), or in the sequencer tracks.

You can also use the "Add Marker" feature (or its keyboard shortcut "**T**") in the "Edit" menu, which places a mark at the start of the selection. During replay, the "**T**" shortcut allows you to set markers "on the fly".

There is a box that you can tick that toggles between hide/show markers in the "**Sonogram Display**" window.

> **Note:** If you have "hidden" the markers, they will not be saved along with the corresponding SDIF file (the file will be empty).

The file containing the markers can be saved under the default name "mysound.mrk.sdif" (in SDIF format) using the "Save Fundamental As. . . " submenu in the "Save Analysis. . . " menu in "File" (see section 29).

These files can be opened, modified and exploited in other applications.

17.2 Automatic Markers

Three types of markers can be generated automatically by analysis:

- Transient attack detection markers, carried out by the new "Transient Detection" algorithm in SuperVP (it is also used for stretching).
- Markers known as "Spectral Flow Markers" (they detect rapid changes in the signal energy. This algorithm is present from AudioSculpt ver 1.2 and greater).There are two options:
  - **"Spectral différencing (positive)"** which detects rapid energy increase in the signal (and which may be linked to a note start);
  - **"Spectral différencing (negative)"** which detects rapid energy decrease in the signal (and which may be linked to a note end).

Parameter choice and settings are carried out in the "**Marker Parameter**" panel that is opened by the "Generate Markers" command in the "Analysis" menu.

In addition to the "FFT Analysis Setting", the Transient Detection Settings must be made. Please see section 27.

Note: By default, the "Verbose Output" box is unchecked. If you check this option (for advanced users: flag -v), all the information provided by SuperVP, except for the command line, will be displayed in the SuperVP console. The information will only be visible if the console is folded out (see section 24.1).

After analysis, the marks will be visible on the sonogram zone.

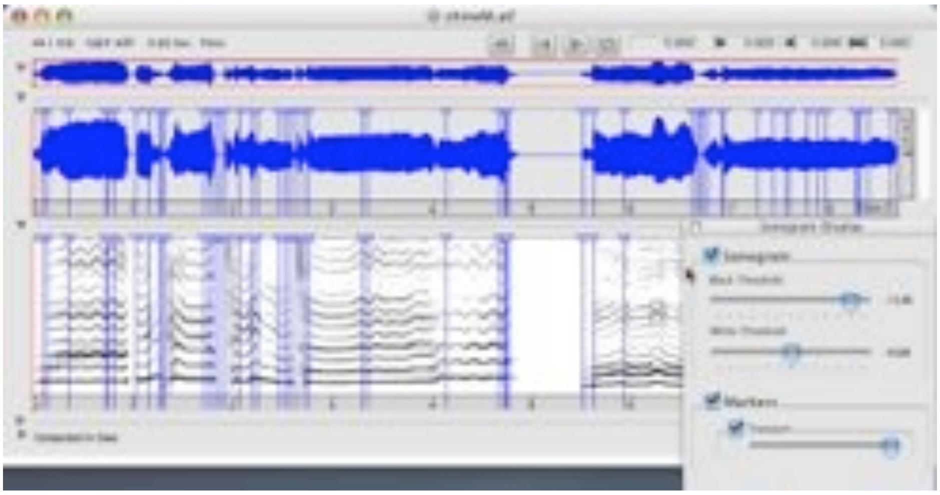

A horizontal slider is displayed in the **Sonogram Display** window (along with the check-box to hide/show the markers that are generated). The slider allows you to set the marker display threshold according to their value.

Note: There is a slider for each of the three marker types.

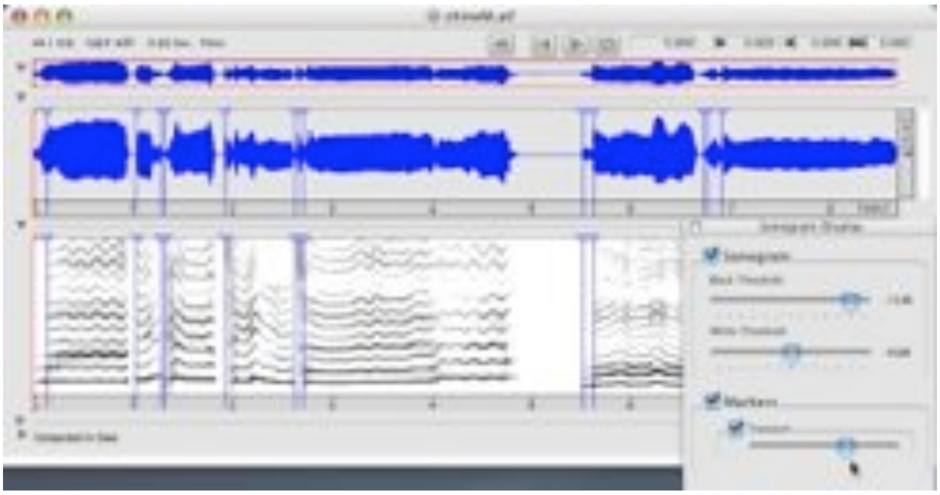

Note: If you have "hidden" the markers, they will not be saved along with the corresponding SDIF file (the file will be empty). If you manipulated the slider, the hidden markers will not be saved in the SDIF file (please see section 29).

These markers can be moved or modified manually via the **Inspector** window. The selection is carried out by means of one of the upper triangles. Multiple selections are carried out in the usual way by holding down the **Shift** key. In this case, only the manual mode will have any effect on them.

Note: The **Inspector** window only acts on the last element selected (please see section 12).

17.3 The Chord Sequence Analysis. . . markers

These are of a different type, since they come in pairs: a start marker and an end marker (please see section 16.5). They are set manually: first, you must select a region. Then, use the Add Chord Seq Markers command (keyboard shortcut "**Q**") in the Edit menu to define the marker pair (start marker and end marker). They are colored red.

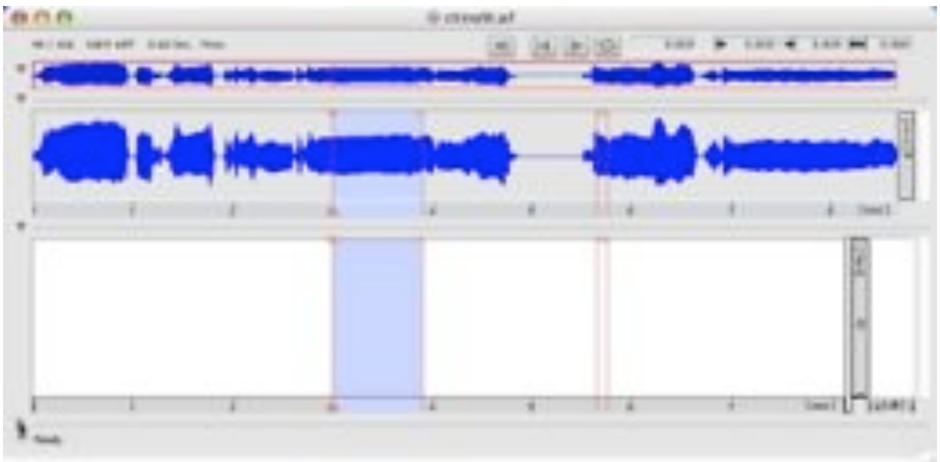

17.4 Marker Colors

Manual markers "Hand Added Markers": red.

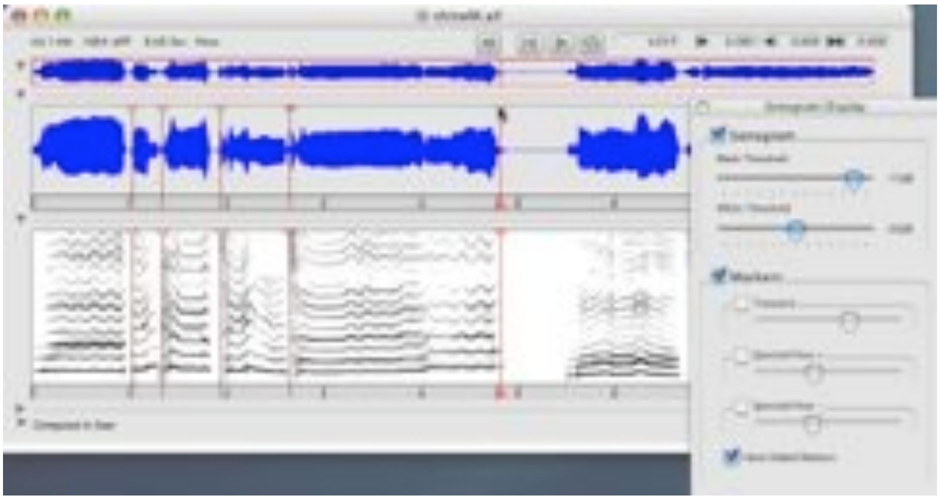

"Transient Markers" : blue

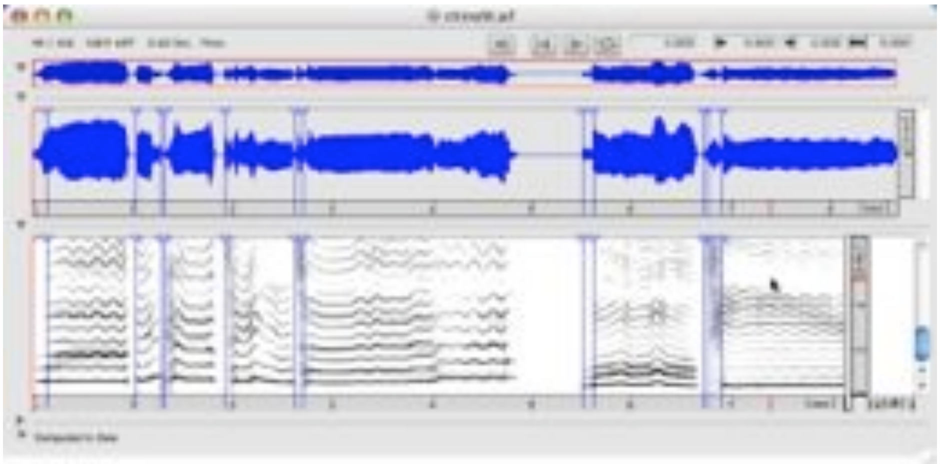

"Positive Spectral Flow Markers" : green

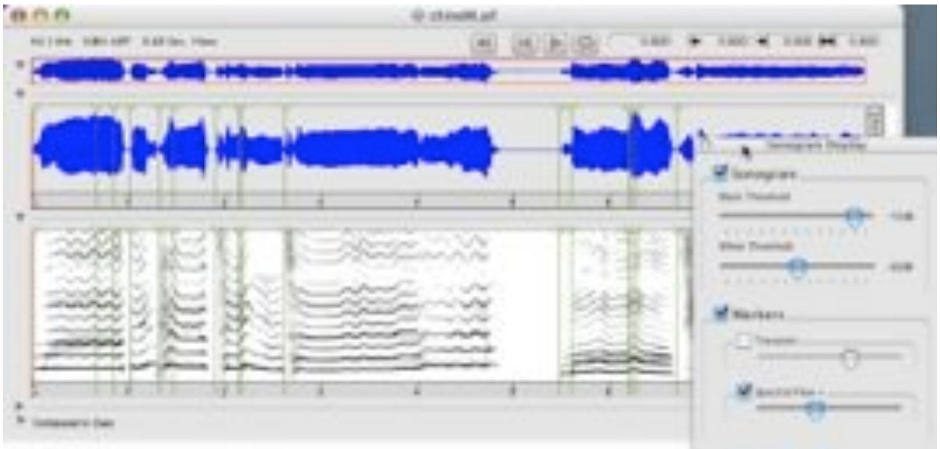

"Negative Spectral Flow Markers" : orange

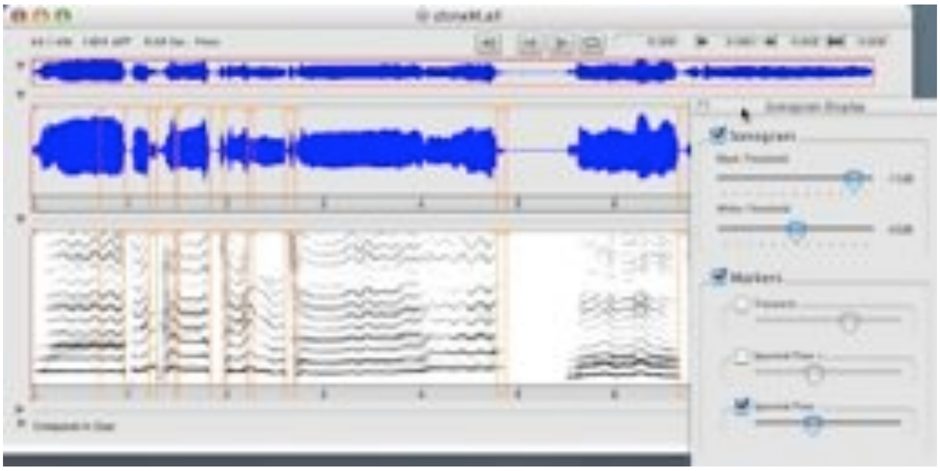

17.5 What all markers have in common

17.5.1 Marker selection

Selection is carried out using any tool (except, of course, the marker tool on one of the upper triangles) . Multiple selection is carried out in the usual way using the **Shift** key.

17.5.2 Moving (modifying) the markers

Markers can be moved manually with any tool (except the tuning fork and the marker tool) via the Inspector window. If several markers are selected, only manual moving can be used.

Note: The Inspector window only affects the last element that was selected (please see section 12).

17.5.3 Deleting Markers

To delete a marker (or several markers in the case of a multi-selection), use the "Delete Selected Markers. . . " command in the Edit menu. All marker types, if selected, will be deleted.

The "Delete Markers" command in the Edit menu opens a control panel allowing you to choose the marker type(s) to be deleted.

17.5.4 Saving your markers

Markers of all types are saved in the same SDIF file, exactly as they appear onscreen. However you can save files under different names.

These files (default name "mysound.mrk.sdif) can be opened, modified and exploited in other applications.

Saving is done as usual via the "Save Markers As. . . " command, in "Save Analysis" in the "File" menu (see section 29).

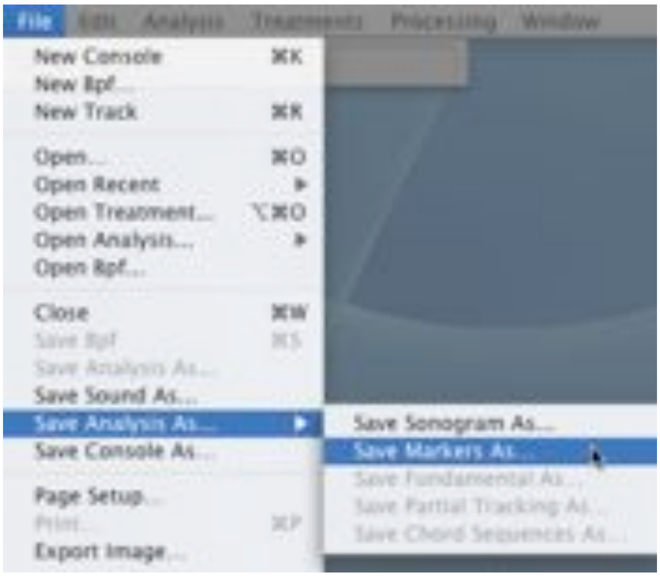

If you close a sound file or try to quit AudioSculpt without saving, the application will ask for confirmation, allowing you to decide whether to save.

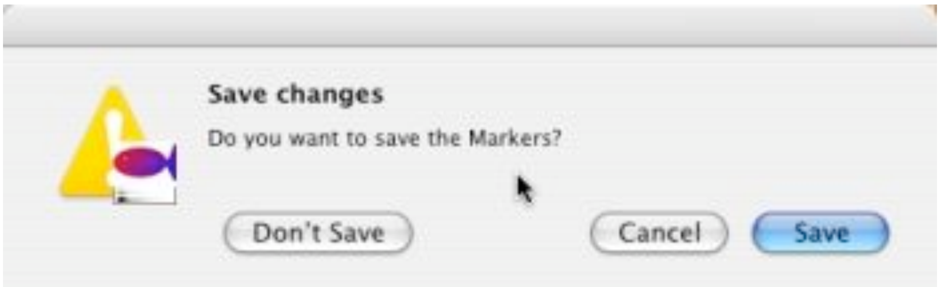

You can name each of them using the "Label" field in the Inspector toolbox.

> **Note:** The "value" field is available for purposes of naming manually placed markers, as well as the "Chord Seq" markers.

17.5.5 Opening Saved Marker Files

While a sound file is open, you can open a previously saved SDIF file, containing the corresponding markers. The "Open Analysis. . . " command contains "Open Markers. . . ". A standard dialog box will open, displaying by default the contents of the "Markers" folder. However you can of course use it to navigate through the hard drive tree and select a .mrk.sdif file.

> **Note:** If a sound file is open, and you request that an analysis file be opened concerning a different sound, AudioSculpt will ask for confirmation. By default, it will always start by searching and opening the sound file that corresponds to the analysis file.

Note: If no sound file is open, and you request that an analysis file be opened, AudioSculpt will open the corresponding sound file.

18 Spectrum and tuning

When you play back all or part of a sound whose sonogram is displayed, the spectrum display (in red) follows the cursor as it progresses.

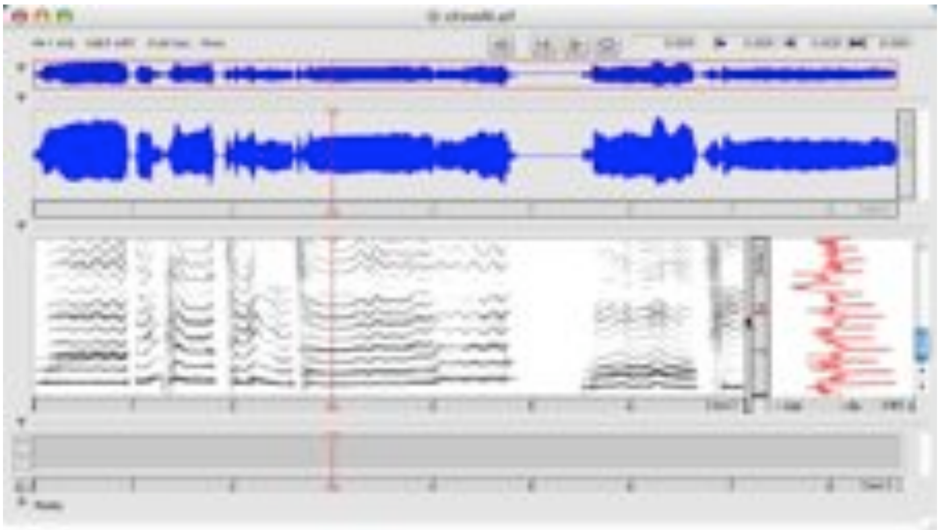

If you place the tuning tool in the sonogram, the spectrum for that point in the sound will turn blue. The active part of the tuning tool is the small red ball at the bottom.

The frequency, amplitude and time position are shown in the right hand part of the toolbox (see section 11.1).

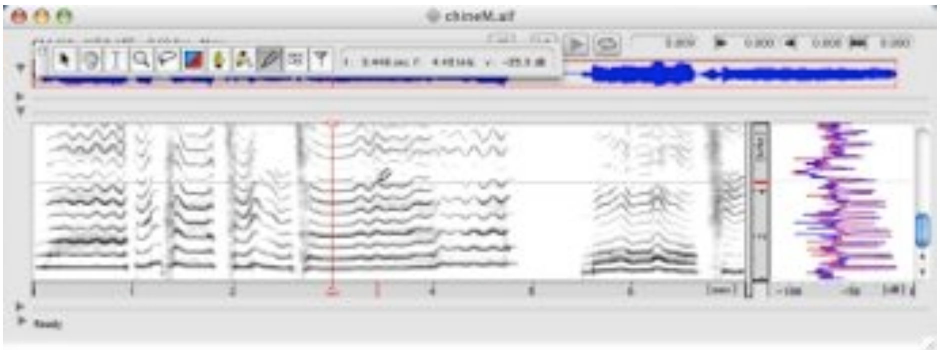

This allows you to compare different parts of the spectrum (red at the cursor position and blue at the tuning position).

Selecting the cursor position is easy: either move it manually, or type in a time value in the first editable field at the top of the window (see section 5.2).

A sinusoidal sound is played at the frequency and amplitude corresponding to the point where you clicked on the sonogram. Frequency, amplitude and time position are shown in the Inspector (see section 12).

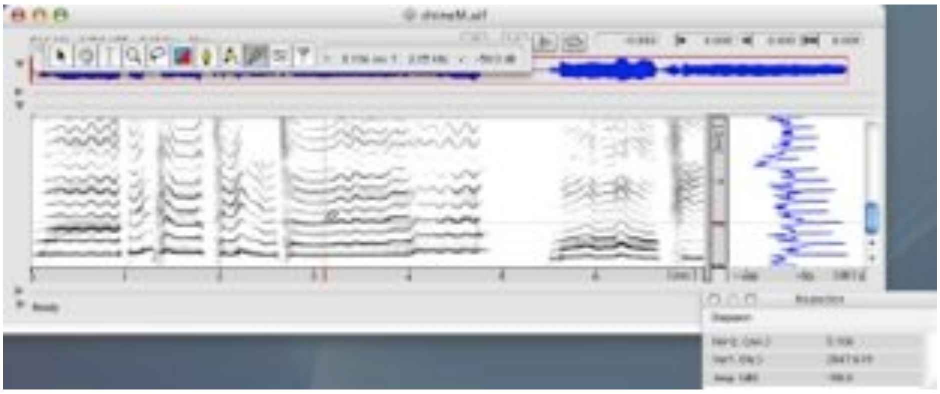

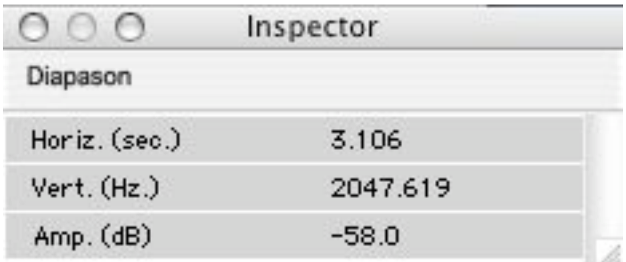

Holding down the **Shift** key allows you to limit dragging to the horizontal or vertical planes. Naturally you can enlarge the zone reserved for the spectrogram (see section 5.2).

You can also zoom using the magnifying tool or the ruler (sections 10.1 and 10.2).

19 Setting up and defining Treatments

19.1 Overview

Setting up and defining Treatments Overview If no selection is made, a treatment is applied to the entire sound by default (otherwise, only to the selected portion).

Two types of treatment are available:

19.1.1 "Constant Parameter" treatments:

Parameters for this treatment type are defined only by constant parameters.

- **Surfaces** : filtering of rectangular or free-hand areas, [or] in pencil strokes.

  Surfaces are only accessible via the Tools toolbox (see section 11).

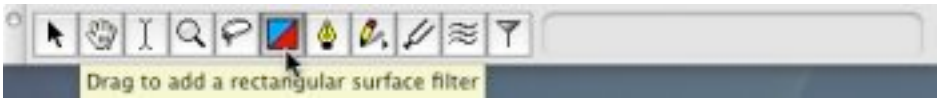

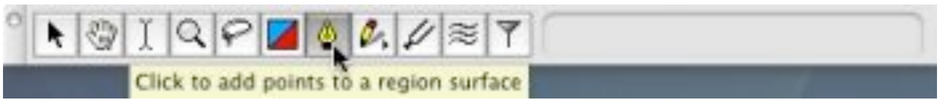

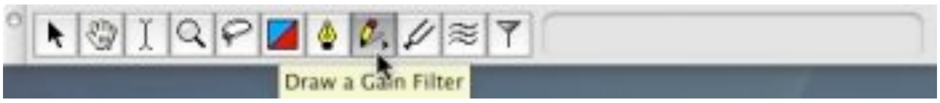

The following Treatments are only accessible via "Treatments" (first part of the menu, please see section 31.4):

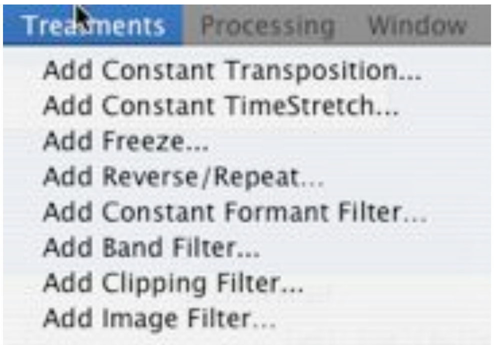

- **Constant Transposition** : constant "Transposition" with or without time correction.
- **Constant TimeStretch** : constant coefficient "Stretching/compression".
- **Freeze** : "freezing" the sound.
- **Reverse/Repeat** : moving back and forth in the sound.
- **Constant Formant Filter** : self explanatory
- **Band Filter** : Bandpass filtering, either "through" or "reject".
- **Clipping Filter** : self explanatory.
- **Image Filter** : uses an image as filter.

19.1.2 BPF processing

This processing requires Bpf definition of certain parameters. A Bpf (Breakpoint Function) is a curve discretely defined by a set of points linked by straight line segments.

This processing is only accessible via the "Treatments" menu (second part of the menu, see section 31.4):

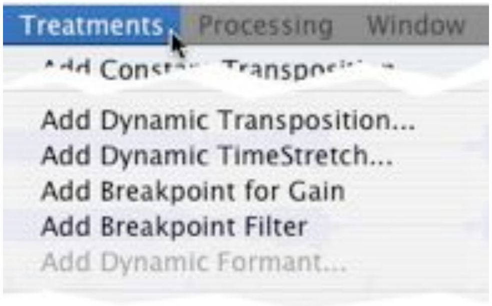

- **Dynamic Transposition** : Time variable transposition.
- **Dynamic TimeStretch** : Variable coefficient stretching/compression
- **Breakpoint for Gain** : Bpf gain modification.
- **Breakpoint Filter** : Bpf filtering.
- **Dynamic Formant** : Formant variable filtering. At the present time, not yet available.

Note: You do not have to display the sonogram in order to apply a treatment.

19.1.3 Treatment display

All processing generates a rectangle known as a "track element" in a sequencer track (see section 21), showing the time position, the duration and an icon of the processing type.

The surfaces are displayed in the sonogram, whether or not the sonogram is currently displayed.

The other processes, if selected, display a rectangle in the sonogram zone to indicate their position. This allows you to place them in the sonogram with precision.

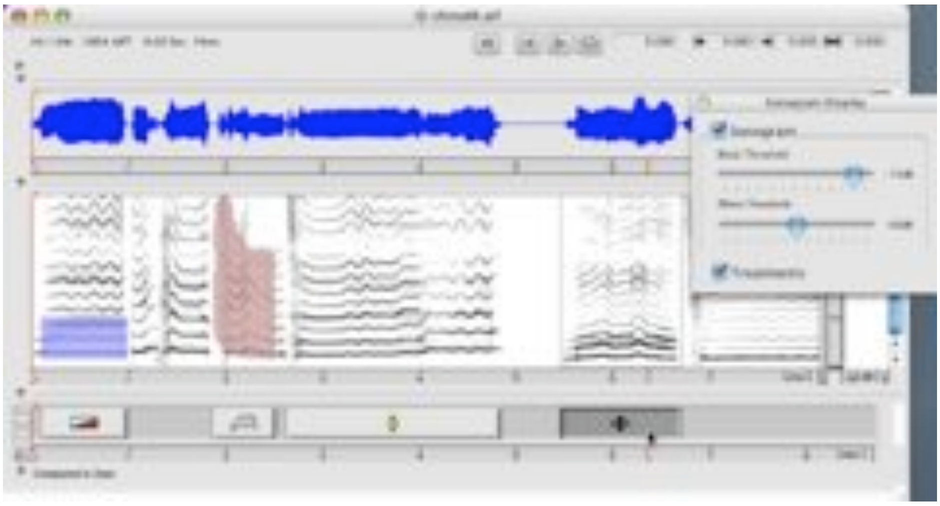

This display can be deactivated on the sonogram by unchecking the "Treatments" box. By default it is checked in the Sonogram Display floating palette. This may be extremely useful for examining the sonogram in greater detail without having to delete the treatments.

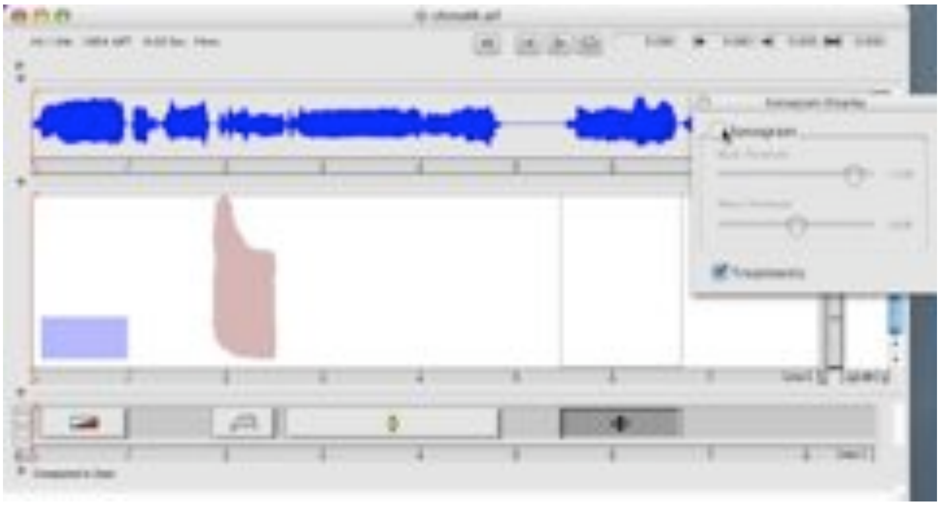

19.2 Surface filtering

Surface filtering :

- Renctangular surfaces
- Free surfaces
- Pencil

Surface filtering is used for modifying surface gain defined in the time/frequency plane. It modifies the amplitude of the surface spectral components. The gain value defines how much this part of the sonogram is amplified or attenuated.

A positive gain value means that the region will be amplified, and a negative value that it will be attenuated.

19.2.1 Renctangular surfaces

Using the rectangle tool, draw a surface on the sonogram. This turns gray, and also appears on the sequencer track (Zone 4) as a track element rectangle bearing the corresponding process icon.

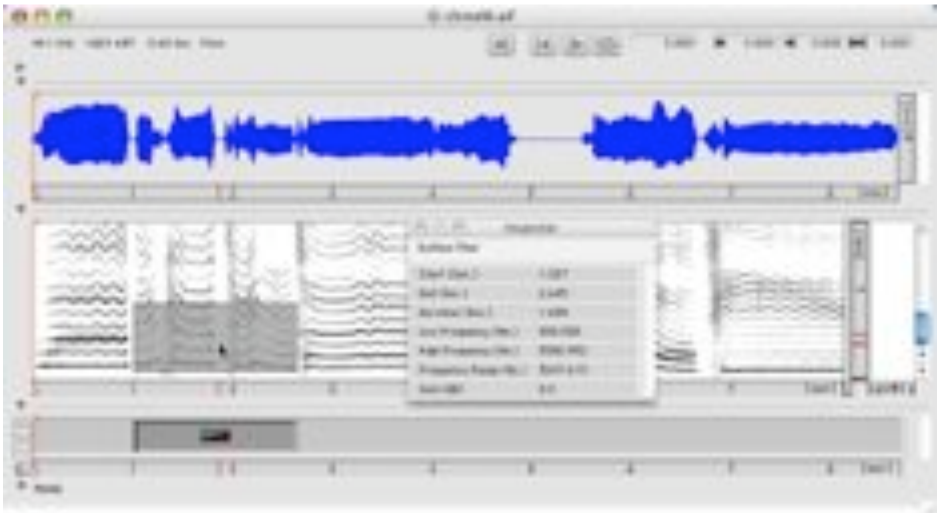

Give it a positive value (in dB) as follows:

First method: Double click on the track rectangle (track element) with any tool (except the marker) to open the "Edit Gain" dialog box. Here you type in the desired value and validate with "OK".

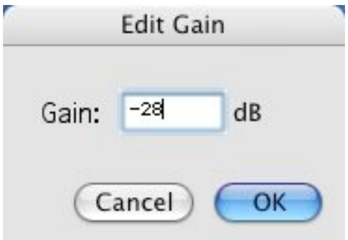

> **Note:** The value limits are -116 dB and +116 dB (if you go beyond these limits, a value of -116 dB or + 116 dB will be taken into account).

Second method: Place one of the following tools: arrow, grab tool, selection cursor or rectangle on the previously drawn surface in the sonogram. Hold down the **Ctrl** key as well as the mouse button and drag to scroll through the values. The value limits are -116 dB and

+116 dB.

Third method: In the Inspector window, type in value in the "Gain (dB)" editable field and press OK.

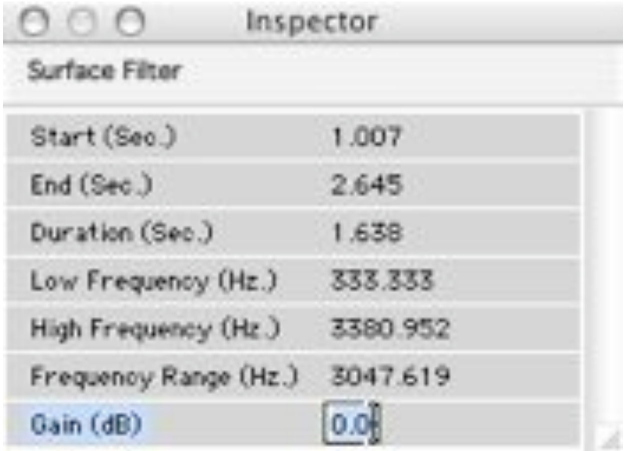

Fourth method: Select "Change Gain. . . " in the Edit menu. The Edit Gain dialog box will open, allowing you to type in the desired value. Then press OK.

The last mentioned method is especially useful for editing the gain of a set of selected processes.

The surface color depends on gain value: greater or less intensity of blue for positive values and greater or less intensity of red for negative values.

You can move the surface that was drawn on the sonogram using the arrow or the grab tool: the track rectangle (track element) will follow. Press the **Shift** key to limit dragging to the vertical or horizontal plane. You can move the track element in time using any tool except the marker tool.

Press on the **Option** (**Alt**) key to duplicate the surface and place the copy at any desired point. Duplicating on the track means you choose only the time position of the copy, but duplicating in the sonogram means you choose both the time position AND the frequency of the copy.

To resize your surface, place the end of the arrow tool on a surface edge. It will change into two little black triangles. Then click to vertically or horizontally enlarge or reduce. The same operation on the track using any tool (except the marker tool) allows you to carry out time changes (see section 11.2.3).

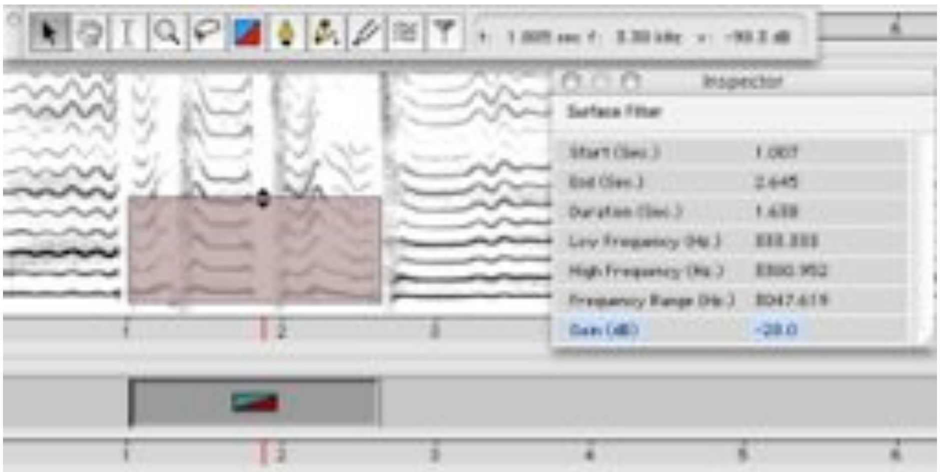

The Inspector window shows the position, the sizes and the gain for the selected surface.

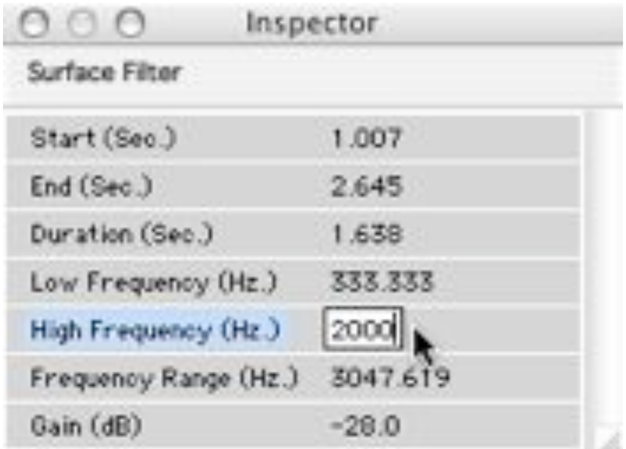

This window allows you to modify any parameter (see section 12).

To delete one or several surfaces, just select the surface(s) with any tool (except the marker tool) on the track, or with the arrow tool in the sonogram, in the usual way (holding down the **Shift** key). Then press the **backspace** key or select "Clear" in the Edit menu. The lasso tool allows you to select several surfaces in the sonogram. You can also operate on the track elements. Please see sections 11.7 and 19.2.4.3.

Drag and drop, used on one or several selected surfaces (or on their corresponding track elements) on the desktop allows you to save all the characteristics of those surfaces (horizontal

and vertical positions as well as gain values)). The treatment will now be called "extract of AudioScultp.textClipping". If you wish to re-use exactly this treatment on another sound, just drag and drop "extract of AudioScultp.textClipping" onto the sonogram (see section 5.2). You can give these extracts any name you wish.

Note: These extracts are automatically numbered, incrementally.

To change the gain of several selected surfaces, use one of the methods described in section 19.2.4.7.

The "Replicate in Frequency. . . " and "Replicate in Time. . . " functions in the Edit menu as well as "Invert" in the "Treatment" menu are active for one or more selected surfaces (see section 19.2.4).

Important: If one or more surfaces overlap, it is better to use the "Maximum" mode in the "Filter Superposition" mode in the "Processing Parameters" panel which is called up by the "Process Treatments. . . " command (see section 25.1.2). This setting is also available via the Preferences filter. . . tab.

19.2.2 Free surface

Using the free-form surface, you can draw 2 surface types in the sonogram: polygons or freehand (see section 11.9).

Draw a polygon by placing the tool at the first desired point, click once, then move the tool to the second desired point (you'll be dragging a segment), click once and so on. Double click to close (to complete) the polygon.

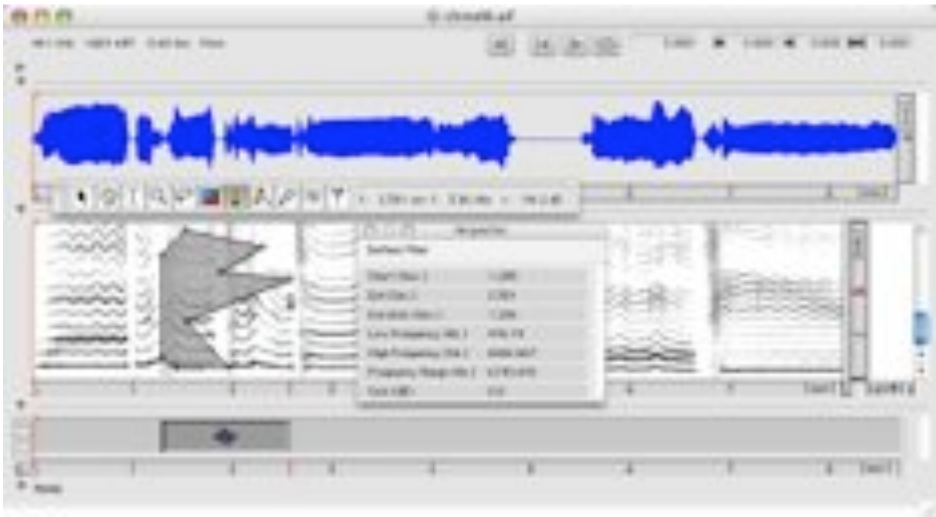

You can move a summit (a crest) using the arrow or the grab tool: the tool will change into a red cross when it passes over a crest. Just click on the crest and move it.

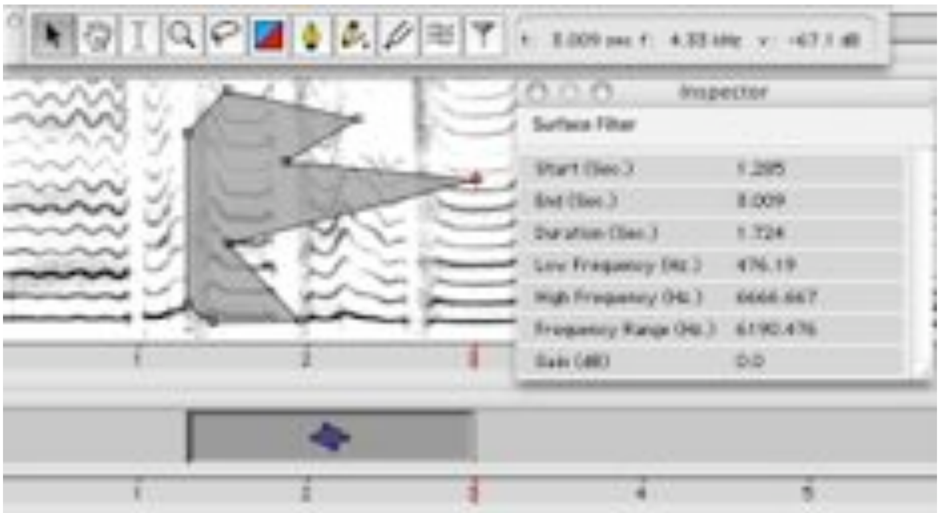

You can create a new summit (a crest) using the arrow or the grab tool: the tool will change into a black cross when it passes over a slope. Just click on the desired spot and move it.

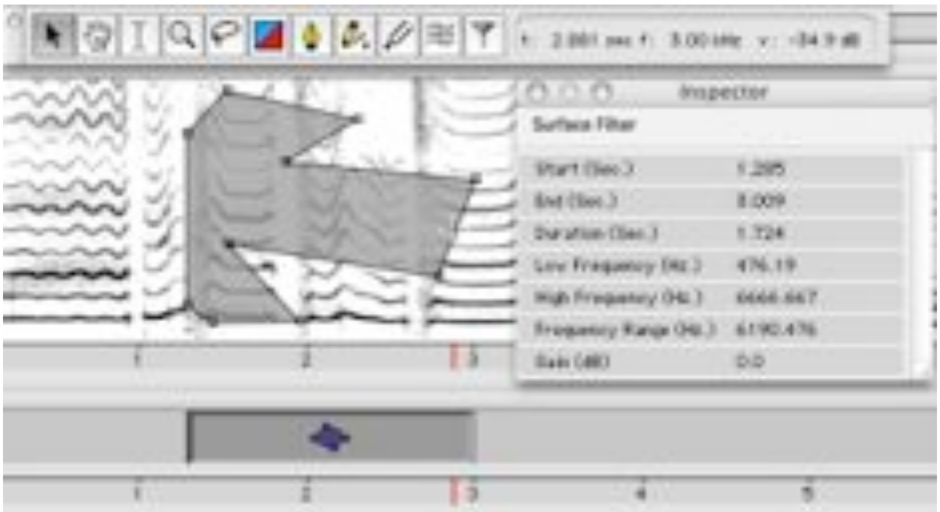

Create a free form surface by placing the tool at the desired point, then hold down the mouse button and draw the desired outline. The surface will close (complete itself) automatically when you release the mouse button.

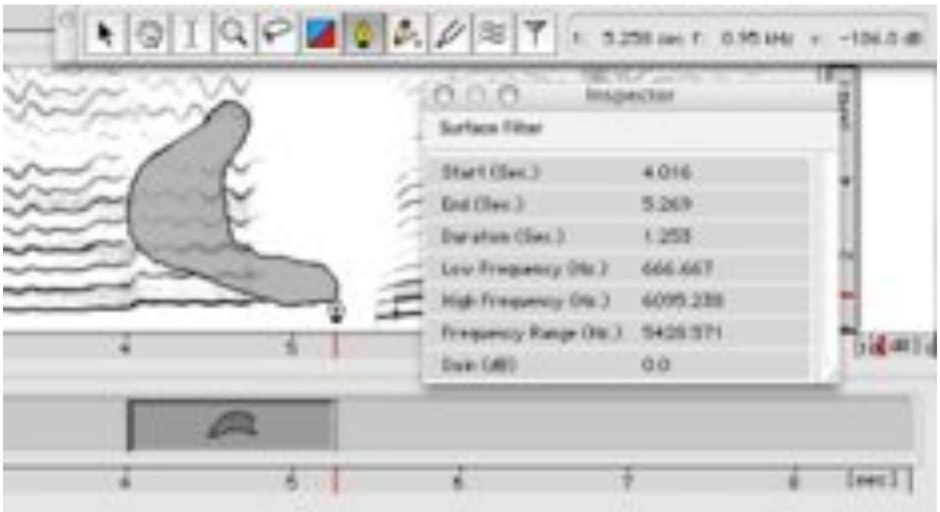

The surfaces are grayed out and also appear on a track (Zone 4) as track rectangles (track elements) of the same duration as each of the corresponding surfaces, each with its own icon.

To give a value in dB to a surface, just use one of the described methods used for rectangular surfaces.

Moving, copying, deleting and saving surfaces via "Drag and Drop" on the desktop is carried out in the same way as for rectangular surfaces.

To resize any of these surfaces along the time axis, just place the end of any tool on a surface edge. It will change into 2 little black triangles. Then click to vertically or horizontally enlarge or reduce.

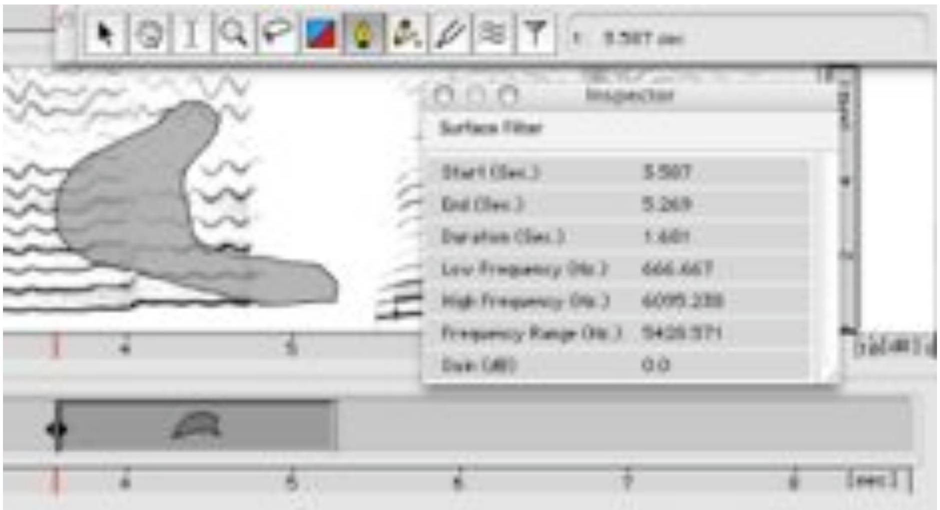

You cannot carry out operations on the surface itself. However, you can modify the highest and lowest frequencies respectively and the frequency width by changing the corresponding parameters in the Inspector window (as well as the gain, the duration and so on).

To change the gain of several selected surfaces, use one of the methods described in

19.2.4.7.

The "Replicate in Frequency. . . " and "Replicate in Time. . . " functions in the Edit menu as well as "Invert" in the "Treatment" menu are active for one or more selected surfaces (see section 19.2.4).

> **Important:** If one or more surfaces overlap, it is better to use the "Maximum" mode in the "Filter Superposition" mode in the "Processing Parameters" panel which is called up by the "Process Treatments. . . " command (see section 25.1.2). This setting is also available via the Preferences filter. . . tab.

19.2.3 The pencil tool

Display of the sonogram is more or less essential for this process.In the Tools toolbox, select the pencil tool as well as the desired line thickness (if the default thickness is not suitable):

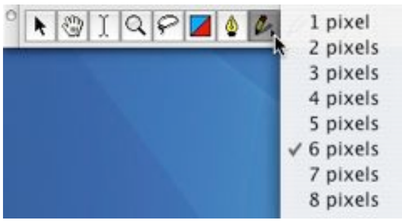

The default line thickness is 6 pixels, and default attenuation is -50 dB.

> **Note:** The default values can be modified. Just click the Pencil Tool tab in the Preferences menu.

> **Note:** To draw straight horizontal or vertical lines, just hold down the **Shift** key.

> **Note:** Vertical lines are not taken into account if they are not thick enough. The [minimum] value depends on the step value (see section 25.1.1).

> **Note:** To reverse the pencil value, just hold down the **Option** (**Alt**) key. The pencil will change signs [positive to negative or vice versa].

In the sonogram, draw a line:

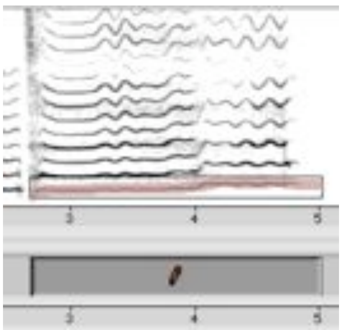

This element behaves EXACTLY like a rectangle surface: horizontal and vertical moving, horizontal and vertical resizing, in the sonogram. The track element will be displayed in a sequencer track. You can move, resize or copy the rectangle (and thus, the process) along the time axis in the usual way on the track.

Moving, copying, deleting and saving surfaces via "Drag and Drop" on the desktop is carried out in the same way as for the other surfaces.

You can also modify the process via the Inspector toolbox: rectangle time and frequency parameters, as well as gain and line thickness ("Line Height" in Hz) are all accessible.

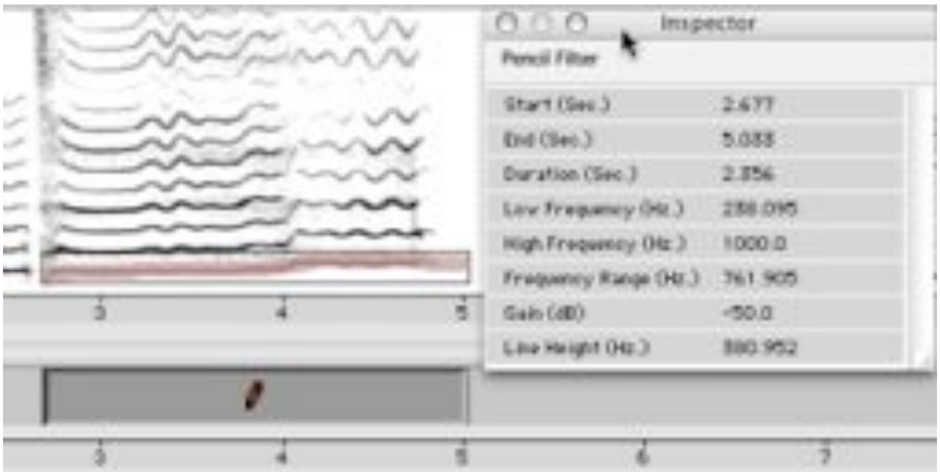

To give another value (in dB) to a pencil line, just use one of the methods described for rectangle surfaces. See section 19.2.1).

To change the gain for several selected pencil lines, use one of the methods described in section 19.2.4.7.

The "Replicate in Frequency. . . " and "Replicate in Time. . . " functions in the Edit menu as well as "Invert" in the "Treatment" menu are active for one or more selected surfaces (see section 19.2.4).

> **Important:** EACH pencil line represents ONE surface: if two lines overlap, the gain values will be accumulated if you are in "Multiply" mode. The "Maximum" mode in the "Filter Superposition" mode in the "Processing Parameters" settings panel is better adapted. You can call this up via the "Process Treatment" command (see section 25.1.2). This setting can also be made via the "Filter" tab in Preferences.

> **Note:** It easily happens that when using the pencil tool, track elements will be superimposed on the same track. In this case, check the "Create Tracks When Needed" box, in the "Treatments" menu (see section 21.2).

19.2.4 What all surface areas have in common (rectangle, free form and pencil)

Please also consult section 19.5: "What all processing treatments have in common."

19.2.4.1 Reminder

Moving, copying, deleting and saving via "Drag and Drop" can be applied to any number of any surface types.

19.2.4.2 Selecting several surfaces

As described above, you can select several surfaces by clicking on them with any tool (except the marker) on the track, or with the pointer (on the sonogram) while holding down **Shift**. Selecting several surfaces on the sonogram can be carried out in the same way as in Finder, using the pointer tool to draw a rectangle around them. Selecting can also be done with the lasso tool (you don’t have to encircle them totally).

19.2.4.3 Deleting several surfaces

To delete one or several selected surfaces, just press the **Delete** key, or the **Backspace** key, or choose Clear in the Edit menu.

Note that this action applies to any selected treatment.

19.2.4.4 Moving surfaces

You can move (a) selected surface(s) as desired around the sonogram, using the arrow or the grab tool: the track rectangle (the track element) will follow.

Hold down the **Shift** key to limit dragging to the horizontal or vertical plane.

You can carry out the same action, but only in the horizontal plane, using any tool (except the marker tool).

Note that this track element action applies to any selected treatment(s).

19.2.4.5 Surface duplication

Press the **Option** (**Alt**) key with the arrow tool or the grab tool to copy selected surfaces and to place to the copy at any desired point. When copying on the track, you can only choose the time position of the copy; when copying on the sonogram, you can choose the time position AND the frequency of the copy.

Note that this action applies to any selected treatment(s).

19.2.4.6 To resize surfaces

To resize a surface, place the tip of the arrow tool on the edge of a surface. It will change into 2 small black triangles; you can then click on these to stretch or reduce the edge in the horizontal plane. When working with rectangle and pencil line surfaces, you can carry out the same action in the vertical plane.

The same operation on the track using any tool (except the marker tool) allows time changes.

Note: Please note that this action can only be carried out on one surface at a time.

19.2.4.7 Gain

You can apply the same gain to several surfaces (rectangles, free form and pencil lines) simultaneously. Select the desired elements as described above, then use one of the following methods:

- First method: Double click on a track element with any tool (except the marker tool), and an Edit Gain dialog box (or an Edit Pencil dialog box if the last element selected was a pencil line) will open. Type in the desired value and validate with "OK".

  Double click on one of the surfaces that were drawn on the sonogram using the arrow tool (or the grab, select or rectangle tools) to produce the same effect.

- Second method: Place one of the following tools on a surface area on the sonogram: arrow, grab, selection or rectangle tool. Then press the **Ctrl** key while holding down the mouse button and moving the mouse. This will allow you to scroll through a succession of values.
- Third method: Choose the "Change Gain. . . " command in the Edit menu. The Edit Gain dialog box will open. Type in the desired value and press OK.

  Note: The value limits are -116 dB and +116 dB (if you go beyond these limits, a value of -116 dB or + 116 dB will be taken into account).

The surface color depends on gain value: greater or less intensity of blue for positive values and greater or less intensity of red for negative values.

19.2.4.8 Overlapping surfaces

Note that if two or more surfaces overlap, the gain values will be accumulated if you are in "Multiply" mode. The "Maximum" mode in the "Filter Superposition" mode in the "Processing Parameters" settings panel is better adapted. You can call this up via the "Process Treatment. . . " command (see section 25.1.2). This setting can also be made via the "Filter" tab in Preferences.

19.2.4.9 "Replicate in Frequency. . . "

This feature is in the Edit menu, and allows you to copy selected surfaces along different parts of the Frequency axis [i.e. from one part of the frequency spectrum to another].

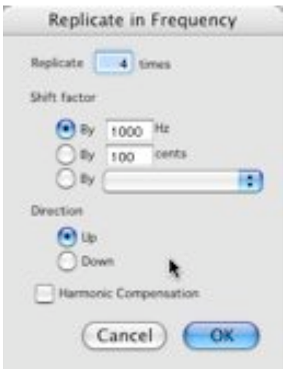

Choose the number of copies and the desired negative or positive frequency difference in Hz or in "cents" (hundredths of a semitone), or in pitch intervals. The "Harmonic Compensation" checkbox allows for automatic adjustment of the harmonics of the frequencies contained in the copy.

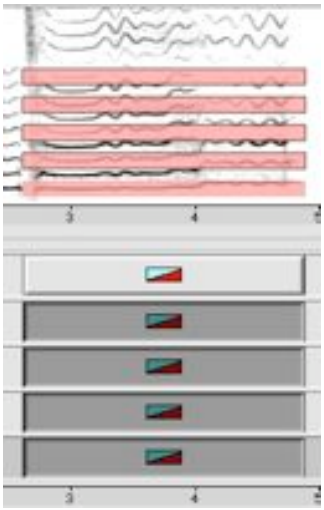

Note that track elements get superimposed if the "Create Tracks When Needed" item in the Treatments menu is unchecked (see section 21.2). Note also that the copies are selected. To change the gain of an individual copy, you must first deselect them (by clicking on the original with the arrow tool or the selection cursor), then by selecting them one at a time using the same tool. Each can be moved horizontally (in fact, they can also be moved vertically, however, this will change their frequency characteristics and defeat your purpose since they would no longer be exact frequency copies).

Note that this feature is only available for surfaces and pencil lines.

19.2.4.10 "Replicate in Time. . . "

This feature is accessed via the Edit menu, and allows you to copy one or more surfaces along the time axis.

Choose the desired number of copies and the desired time-shift to the right (positive value in seconds) or to the left (negative value in seconds).

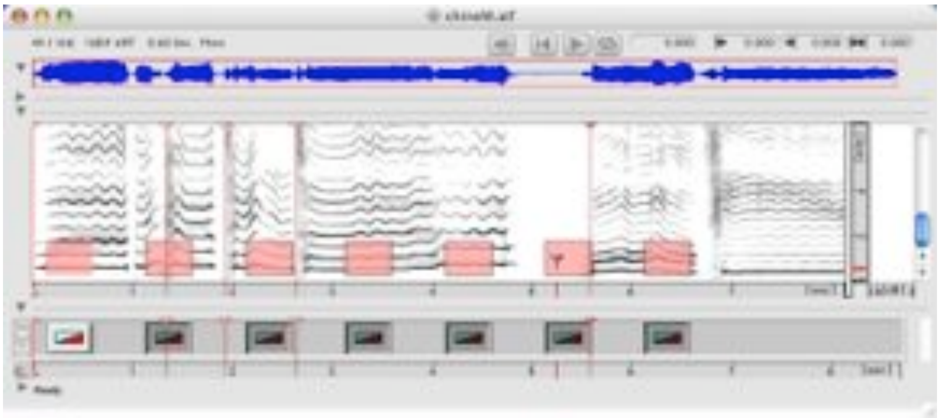

The "Align to Markers" checkbox can only be checked if there are markers. They allow you to place copies on markers located at a time point higher or equal to the desired time-shift.

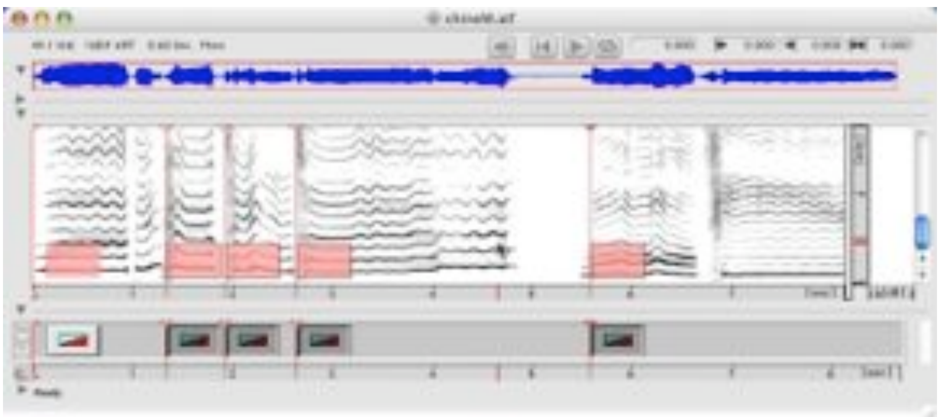

Important: If the requested number of copies is greater than the number of usable markers, the additional copies will be superimposed (stacked) on the last copy.

Note that this feature can be used on any treatment.

19.2.4.11 "Expand To Markers"

This feature is accessed via the Edit menu (keyboard shortcut: **X**) and allows you to stretch the selected surface(s) so as to fit the nearest markers on the left and the right.

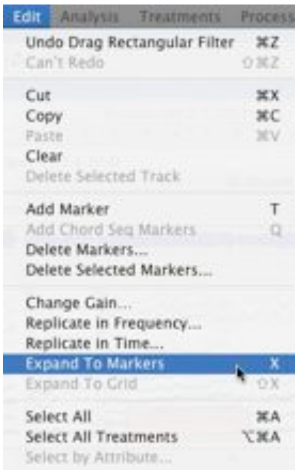

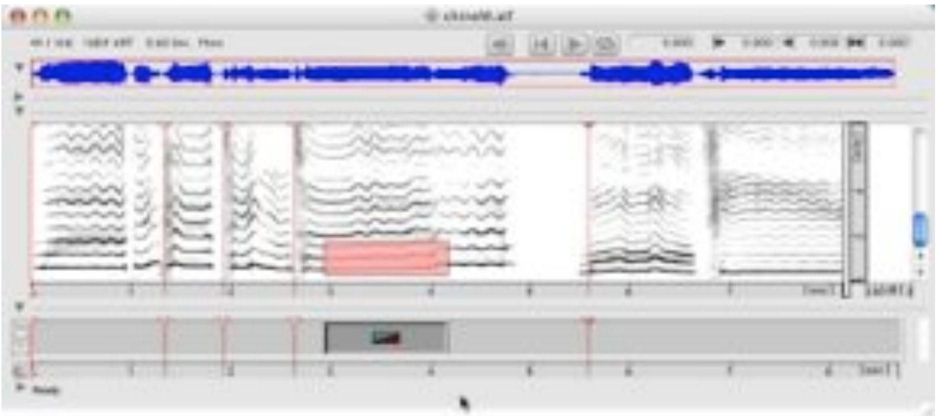

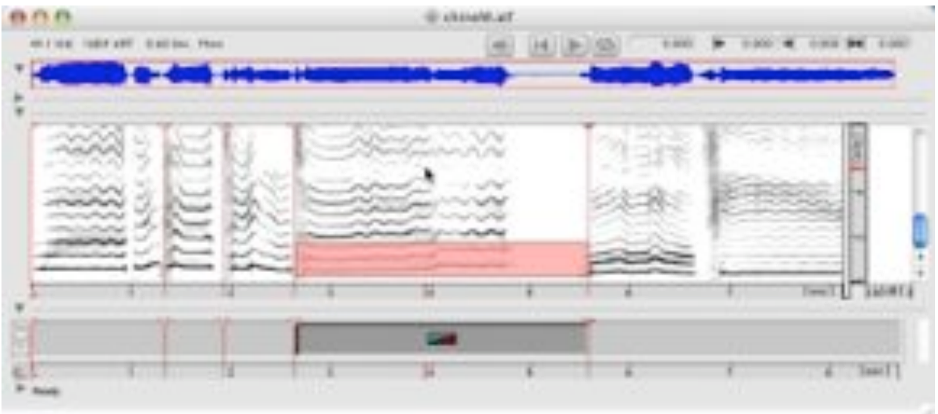

19.2.4.12 "Expand To Grid"

This feature is accessed via the Edit menu (keyboard shortcut: **Shift-X**) and allows you to stretch the selected surface(s) so as to fit the nearest grid references to the left and to the right.

19.2.4.13 "Invert"

This feature is accessed via the Treatments menu and allows you to change the gain sign of the selected surface(s).

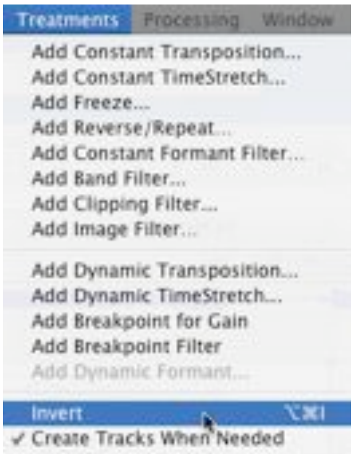

Note that this feature can be applied to any treatment (see below) except the "Clipping Filter". According to where it is used, it will act upon the gain or some other parameter.

19.2.4.14 "Drag and Drop"

Dragging and dropping one or more surfaces (or their corresponding track elements) on the desktop allows you to save all the characteristics of those surfaces (their horizontal and vertical positions as well as their corresponding gain values): the treatment (process) will automatically be named "Extract of AudioSculpt" [AudioSculpt.textClipping extract] or, with the file extension, "Extrait de AudioSculpt.textClipping" [AudioSculpt.textClipping extract]. If you wish to apply this treatment, with identical parameters, to some other sound file, simply drop the icon of the .textClipping file (which can be given any name you desire) onto the sound file (see section 5.2).

Note: these extracts are automatically numbered incrementally. Note that this operation can be applied to any selected treatment(s).

19.3 Other "Constant Parameter" Treatments

If no part of the sound is selected, the treatment will be applied to the entire sound.

The treatments available are in the first part of the Treatments menu:

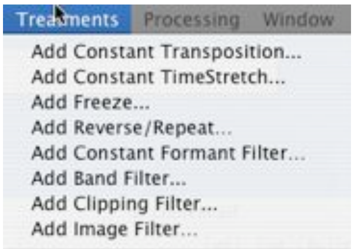

- **Constant Transposition** : applies a constant transposition, i.e. without time correction
- **Constant TimeStretch** : applies a variation in the time domain, defined by a constant coefficient.
- **Freeze**: allows you to "freeze" a sound.
- **Reverse/Repeat** : navigates back and forth along the sound.
- **Constant Formant Filter** : fixed formant type filtering.
- **Band Filter** : Bandpass filtering (negative or positive)
- **Clipping Filter** : "Clipping" type filtering.
- **Image Filter** : "Clipping" type filtering using an image.

19.3.1"Constant Transposition"

In the "Treatments" menu, select "Add Constant Transposition. . . ". The dialog box will open and allow you to set the desired parameters.

First choose "up" or "down" to determine the transposition direction. It can be set in "cents" (one hundredth of a semitone) or in intervals. A checkbox allows you to decide for or against time correction ("Time Correction" is checked by default). If you choose time correction, you

can decide whether or not to keep the spectral envelope ("Preserve Envelope" is checked by default). These options are available in the "Inspector" window.

Transposition is limited to +5 or -5 octaves, a considerable range.

Below that, a grayed out line indicates the compression/stretching amount that would be required in order to obtain the same degree of transposition without time correction.

If you deactivate time correction, the corresponding stretch/compress coefficient field becomes accessible, allowing you to choose the corresponding degree of transposition (the value in "cents" is shown above). The value that is applied will depend on the choice of transposition direction (by having checked "Up" or "Down").

At the same time a rectangle having the same duration will appear on the sequencer track, carrying the icon that corresponds to the treatment (track element).

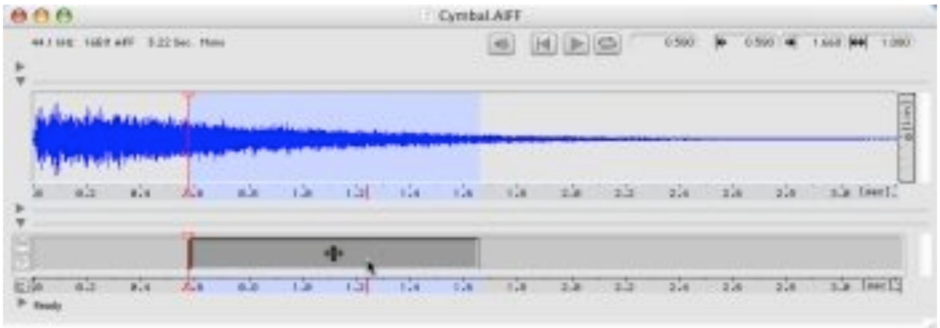

In order to change the transposition, you need to open the dialog box by double clicking on the track element.

You can set transposition mode (with or without time correction) separately for each of the transpositions. You can make this choice for each treatment either in the usual dialog box or in the "Inspector" window (by toggling on/off opposite "Time Correction"). You can make the same choice for "Preserve Envelope" if "Time Correction" has been activated.

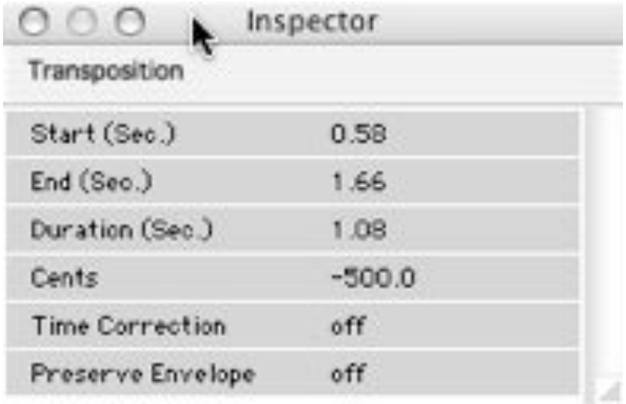

In the control panel that opens when you choose "Process Treatments. . . ", "Process Selection. . . " or "Real-time Processing Settings. . . " ("Processing" menu), the "Preserve Transient" option allows for the detecting of transients in order to improve the transposition (as well as stretching/compression). This option uses predefined parameters. Transient detection is only possible in "Phase Synchronous Processing" mode (see section 25.1.2).

You can move or resize along the time axis, or you can duplicate the rectangle (and thus duplicate the treatment) on the track in the same way as with the "surface" elements (see section 19.5).

You can also modify your treatment via the "Inspector" toolbox (palette), where you can change parameters as desired (see section 12).

The "Replicate in Time" (Edit menu) and "Invert" (Treatments menu) are activated. The "Invert" feature changes the transposition sign (the direction).

19.3.2 "Constant TimeStretching"

In the "Treatments" menu, select "Add Constant TimeStretch. . . ". The dialog box that follows will allow you to choose the stretch/compression coefficient.

The maximum compression ratio is 0,01 and the maximum stretch ratio is 1000 (if you go beyond those limits, the program will automatically correct the value).

At the same time a rectangle having the same duration will appear on the sequencer track, carrying the icon that corresponds to the treatment (track element).

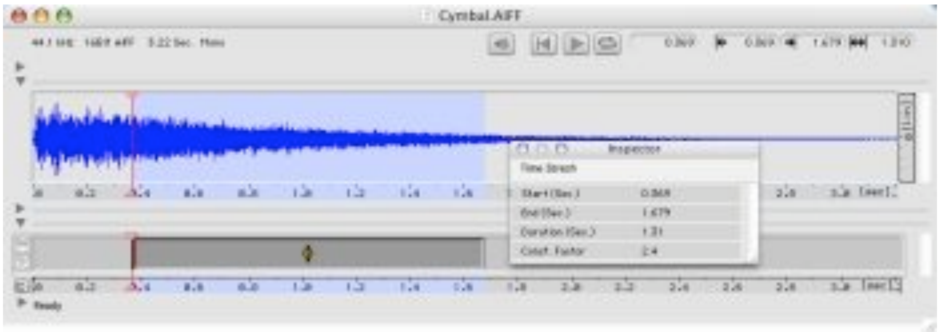

In order to change the coefficient, you need to open the dialog box by double clicking the track element.

In the control panel that opens when you choose "Process Treatments. . . ", "Process Selection. . . " or "Real-time Processing Settings. . . " ("Processing" menu), the "Preserve Transient" option allows for the detecting of transients in order to improve the transposition (as well as stretching/compression). This option uses predefined parameters. Transient detection is only possible in "Phase Synchronous Processing" mode (see section 25.1.2).

You can move or resize along the time axis, or you can duplicate the rectangle (and thus duplicate the treatment) on the track in the same way as with the "surface" elements (see section 19.5).

You can also modify your treatment via the "Inspector" toolbox, where you can change parameters as desired (see section 12).

The "Replicate in Time" (Edit menu) and "Invert" (Treatments menu) are activated.

The "Invert" feature generates the opposite coefficient: y=1/x (e.g. 2 becomes 0,5).

19.3.3"Freeze" : freezing the sound at a given point

This treatment is applied to a given point in the sound. This point (or segment) will be stretched to the "freeze" duration value.

Choose "Add Freeze" in the "Treatments" menu to open the control panel that will allow you to set the freeze duration. The part of the sound that is analyzed for freezing is determined by selection (History Duration). The "freeze" is carried out at the end of the selection [Translator's note: it is not clear whether this means "after the selection has been made" OR "at the end part of the selection"].

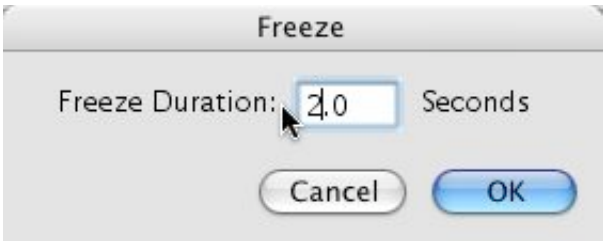

At the same time a rectangle having the same duration will appear on the sequencer track,

carrying the icon that corresponds to the treatment (track element).

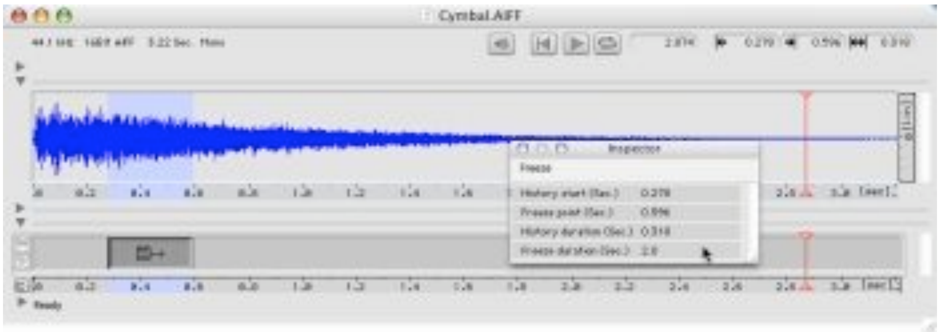

In order to change the treatment, you need to open the "Freeze" settings panel (control panel) by double clicking the track element. The "Inspector" palette allows you to change the parameters: analysis startpoint (History Start), Freeze Point, History Duration and Freeze Duration.

The result:

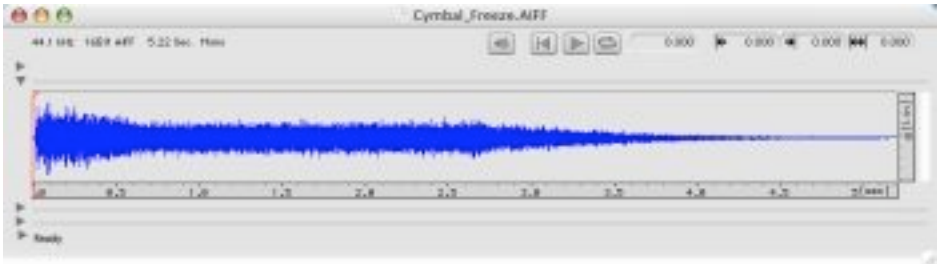

The "Replicate in Time. . . " function is active.

The "Invert" function is not active.

19.3.4 "Reverse/Repeat"

This treatment is applied to all or part of the sound. You can play the sound forwards or backwards, with or without looping. The resulting sound is saved, and can accordingly be longer than the original. SuperVp allows you to apply this treatment seamlessly and without clicks or pops.

In the "Treatments" menu, select the "AddReverse/repeat" item to open the mode setting panel ("Reverse Mode"):

- "Backwards"
- "Back And Forth"
- "Forward"
- "Forth And Back" :. . . and the number of times you want the sound to loop ("Repeat Count").

Note: One pass is outward OR return.

For example: setting the mode "Forth and Back" with a Repeat Count = 1 will cause no changes to the sound, since it will play forward just once, through the selection.

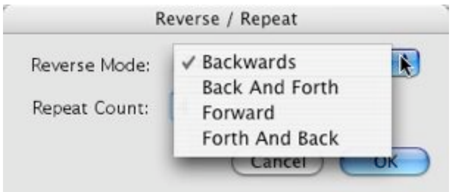

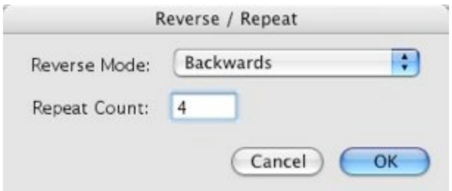

At the same time a rectangle having the same duration will appear on the sequencer track, carrying the icon that corresponds to the treatment (track element).

In order to change the treatment, you need to open the "Reverse/Repeat" settings panel (control panel) by double clicking the track element.

The "Inspector" toolbox allows you to change all the parameters.

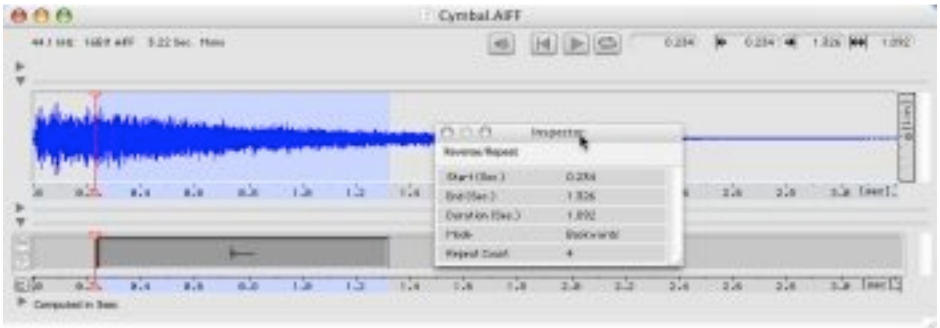

The "Replicate in Time. . . " function is active. The "Invert" function is not active.

19.3.5 "Constant Formant Filter"

A formant type filter is modeled on vocal formants: therefore it is made up of positive peaks.

When you increase formant gain, all other values are increased accordingly. Similarly, when you reduce formant gain, all other values are correspondingly reduced.

In the "Treatments" menu, select " Add Constant Formant Filter. . . ". The "Formant Filter" dialog box will open and allow you to set the desired parameters.

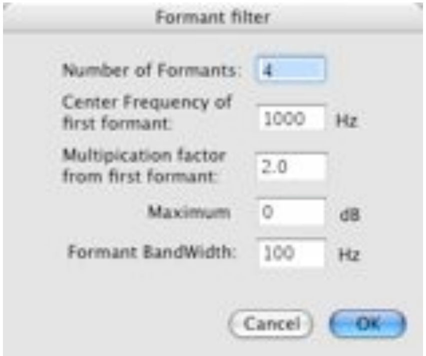

At the same time a rectangle having the same duration will appear on the sequencer track, carrying the icon that corresponds to the treatment (track element).

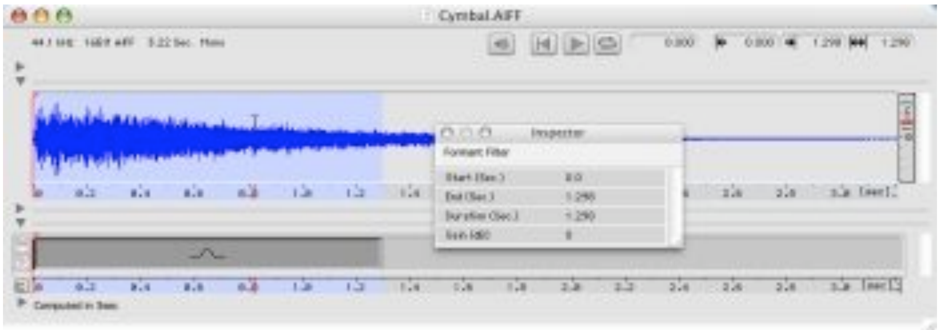

To change the treatment, you need to open the dialog box by double clicking the track element. You can move or resize along the time axis, or you can duplicate the rectangle (and thus duplicate the treatment) on the track in the same way as with the other elements (see section 19.5). You can also change the treatment using the "Inspector" palette. But you only have access to the time position, the duration and the gain.

Only the "Replicate in Time..." function is active (Edit menu). The Invert function is not active.

19.3.6 "Band Filter"

In the "Treatments" menu, select "Add Band Filter"..." A dialog box opens and allows you to set the desired parameters.

At the same time a rectangle having the same duration will appear on the sequencer track, carrying the icon that corresponds to the treatment (track element).

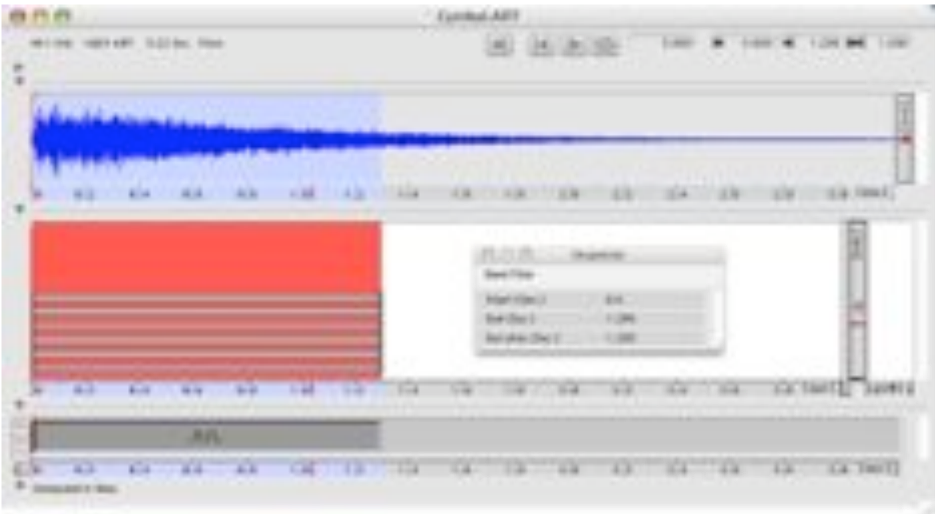

You can edit the bands as desired. You can move a point, using the pointer tool. It will turn into a red cross when it rolls over one of them. Just click on it and move it to the desired place.

You can create a pair of points at a given time spot, using the pointer tool: It will turn into a black cross when it rolls over a band limiting line. Just click on it and move it to the desired place.

To delete a pair of points, just place the pointer tool upon one of them, and it will turn into a red cross. Then you simply click while holding down the "**Option**" (**Alt**) key. Bands may cross (overlap).

If you open the dialog box by clicking on the track element, you can only choose between "Band Pass" and "Band Reject", without editing the bands. You will obtain the same [analogous] result using the "Invert" function.

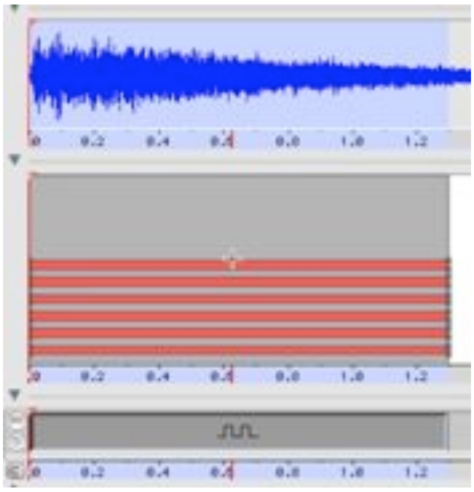

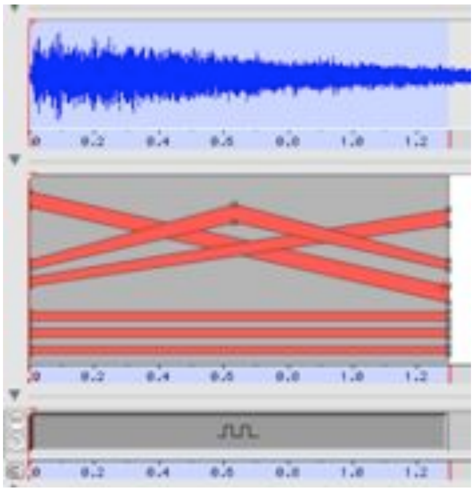

But if you choose "More Options", AudioSculpt will warn you that you are about to lose all previous changes.

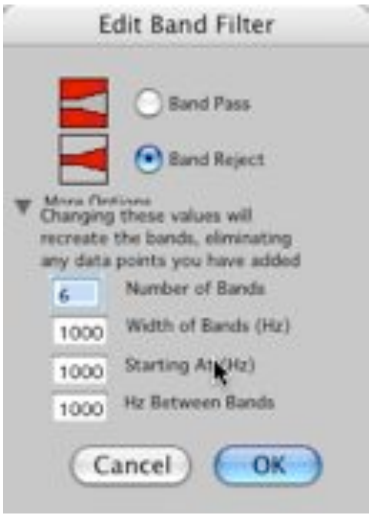

You can move or resize along the time axis, or you can duplicate the rectangle (and thus duplicate the treatment) on the track in the same way as with the other elements (see section 19.5). You can also move the selection vertically on the sonogram using the pointer tool. You can also change the time parameters using the "Inspector" toolbox.

The "Replicate in Time. . . " (Edit menu) and "Invert" (Treatment menu) are activated. The Invert function inverts "Band Pass" and "Band Reject".

19.3.7 "Clipping Filter"

"Clipping" filtering allows you to carry out threshold filtering, which depends on the sonogram's gray scale thresholds: This requires that you first calculate and display the sonogram.

In the Treatments menu, use the Add Clipping Filter. . . item to add a Clipping track element. Simply click on it to open the Edit Clipping Filter settings panel.

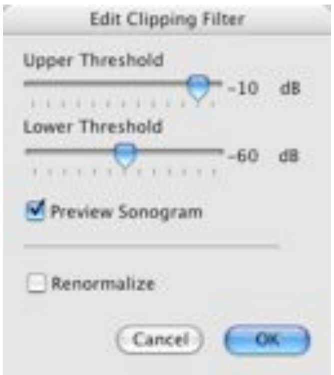

Setting the sonogram display, and thus the filtering display, is done using the "Upper Threshold" and "Lower Threshold" faders. The default settings are: -10 dB for black and -60 dB for white. These are the same values as the sonogram's default display. By checking the "Preview Sonogram" box, you will be able to visually evaluate the kind of change carried out

on the entire sonogram (however, the treatment is only applied to the selected part). The Renormalize checkbox allows you to compensate for signal loss relative to the rest of the sound.

At the same time a rectangle having the same duration will appear on the sequencer track, carrying the icon that corresponds to the treatment (track element).

Original sonagram :

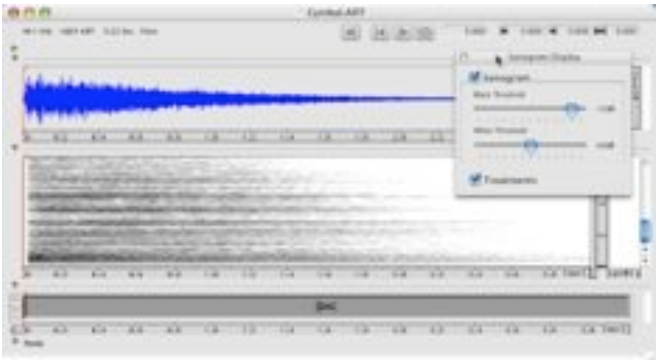

Threshold setting :

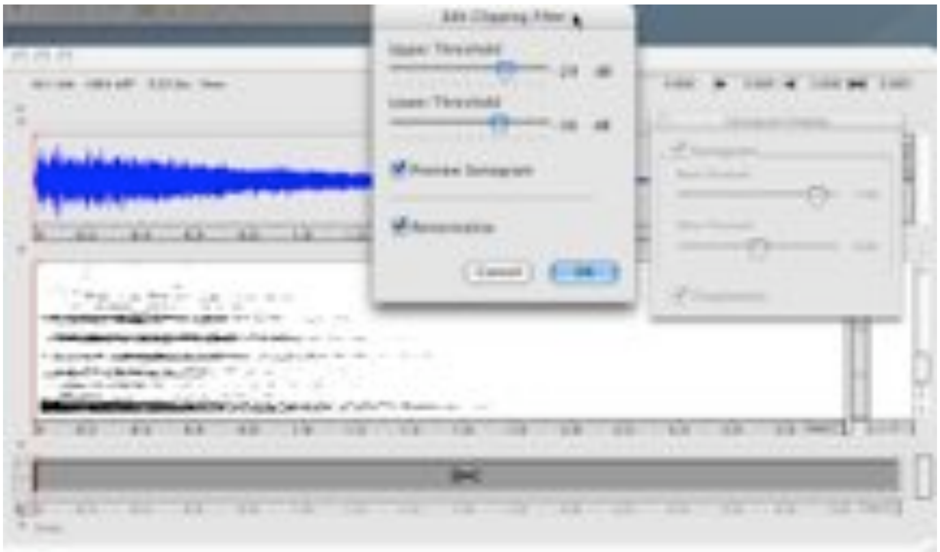

To change the treatment, you need to open the Edit Clipping Filter by double clicking the track element. You can move or resize along the time axis, or you can duplicate the rectangle (and thus duplicate the treatment) on the track in the same way as with the other elements (see section 19.5). You can also edit your treatment using the Inspector toolbox, setting the parameters as desired (see section 12).

The "Replicate in Time. . . " function is active (Edit menu). The Invert function is not active.

19.3.8 "Image Filter"

This is done by placing an image on all or part of the sonogram and using it as a grayscale filter.

In the Treatments menu, the Add Image Filter item will open the standard dialog box that allows you to navigate around the hard disk tree, and choose an image (a photo, a drawing etc) to be imported. This can be in color or in grayscale, in the JPEG, PICT, TIFF, PDF formats. . . any format supported by Quicktime. The selected image will be displayed in grayscale at the desired place on the sonogram.

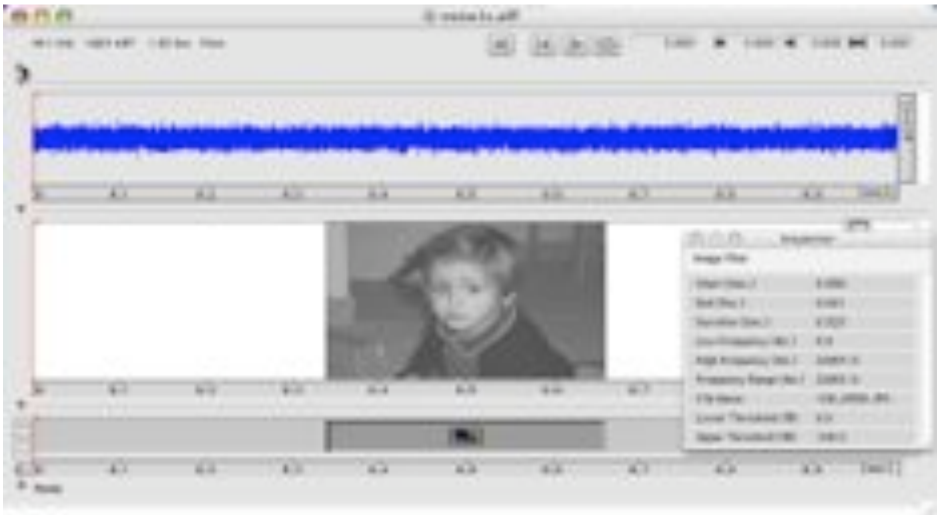

Its shape can be changed, according to the sonogram boundaries. This image has the same properties as a rectangular surface: vertical and horizontal resizing, and moving in any direction around the sonogram (see sections 19.2.4.4 and 19.2.4.6).

At the same time a rectangle having the same duration will appear on the sequencer track, carrying the icon that corresponds to the treatment (track element). Simply click on it to open the Edit Image Filter settings panel.

The name of the selected image appears upper left. At right, the Select button allows you to choose another image. There are the same faders (Upper Threshold and Lower Threshold) as for Clipping, allowing you to change the black and white thresholds .

However, the default settings (expressed in percent) are: 100 % for Upper Threshold and 0 % for Lower Threshold, adhering to the image's dynamics. If the image does not cover the

entire frequency band, frequencies not covered will be deleted (default setting). The "Pass frequencies outside of the image's frequency range" allows you not to delete them.

To change the treatment, you need to open the Edit Image Filter by double clicking the track element. All parameters can be edited in the Inspector toolbox, except choice of image.

The "Replicate in Time. . . " function is active (Edit menu). The Invert function (Treatments menu) is not active. It generates the image's negative.

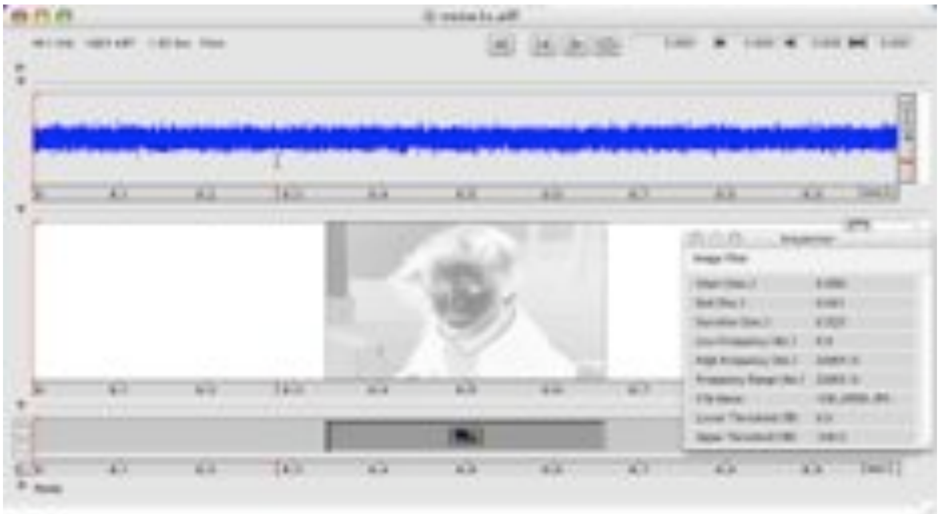

19.4 "Bpf" treatments

If no part of the sound is selected, the defined treatment will be applied to the entire sound.

19.4.1 The Bpf editor

A Bpf (Breakpoint Function) is a curve discretely defined by a series of points linked by straight lines.

The Bpf editor is used by Diphone. Please refer to the Diphone documentation, on the CD Rom: "Diphone_Studio-English.pdf" (section "2.6 Actions on Bpf's"), and its update "Diphone 4r2n MAJ-E.pdf" (section 21 : The independent Bpf editor).

Each Bpf is identified by the name "sound, Bpf type, start time and end time". The window name is useful for identifying it and locating it (via the Windows menu) among the numerous windows that might be open.

To edit a Bpf, open the editor by double clicking on the track element (or on the corresponding sonogram surface). Modifications are carried out as for all Bpf's: using the pencil, the eraser or the "Bpf Tools".

> **Note:** Hold down the **Option** (**Alt**) key to force the pencil tool to draw straight lines (horizontal or slanting). A small line will appear beside the pencil.

If the Bpf resolution is too low, especially for Bpf filtering (Breakpoint Filter), simply change the sample rate using the Resampling item in the Bpf menu.

Don't forget the useful little tool for resizing Bpf windows.

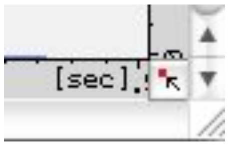

If you select several Bpf treatments (be they identical or different) only the Bpf that corresponds to the track element you double clicked upon will open. Double clicking on the sonogram will open the Bpf that corresponds to the last selected treatment.

19.4.2 "Dynamic Transposition"

In the Treatments menu, the Add Dynamic Transposition... item will open the Bpf editor allowing you to define the transposition coefficient as a time function. This window (and thus, the Bpf) is identified by the name "sound, Bpf type, start time and end time" (in the present case, "0.681.8): Here you can draw and manipulate the Bpf that represents your variable coefficient. The window name is useful for identifying it and locating it among the numerous windows that might be open, thanks to the Windows menu.

The horizontal axis is time, and the vertical axis is transposition in "cents" (hundredths of a semitone) or in semitone intervals.

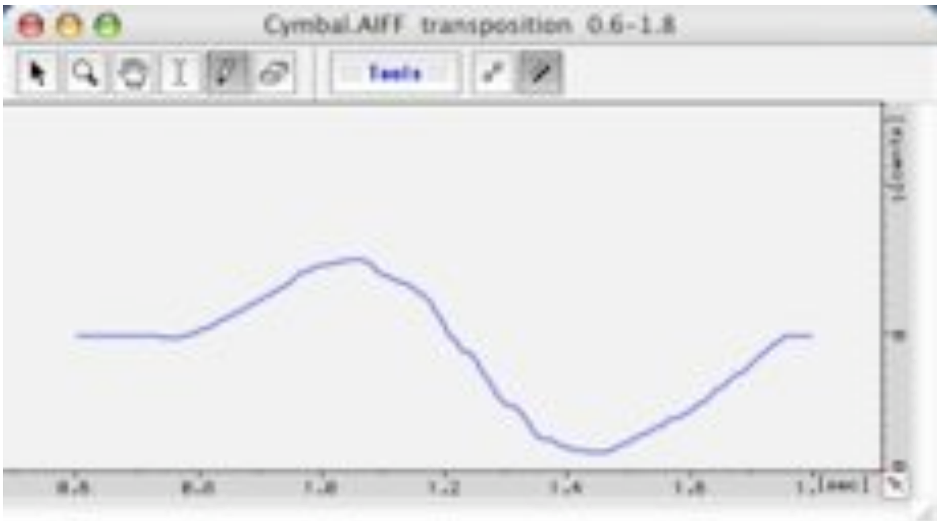

The editor limits transposition to +5 or -5 octaves, a considerable range (to be used with caution!). The corresponding track element will be positioned on the sequencer track.

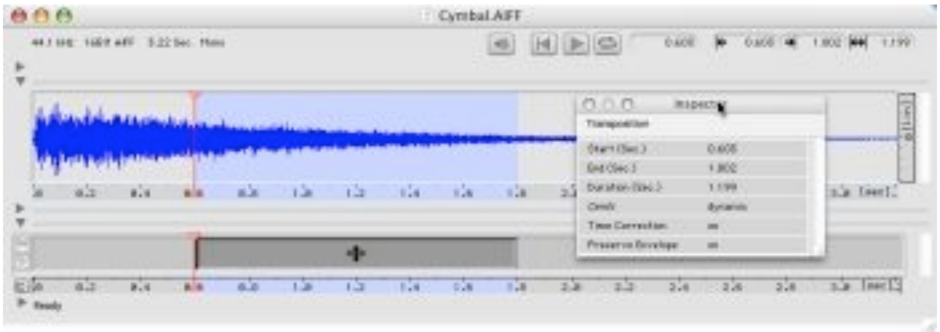

The Inspector toolbox allows access to the following parameters: start, end and duration, as well as time correction and spectral envelope preservation.

Each transposition can be carried out with or without time correction. Use the Inspector window to choose this parameter independently for each treatment (toggle on/off opposite "Time Correction). The same (analogous) choice is possible for "Preserve Envelope", on condition "Time Correction" has been checked.

To edit the Bpf, double click on the track element (or on the corresponding sonogram surface). Modifications are carried out as for all Bpf's. If you select several Bpf treatments (identical or different) only the Bpf that corresponds to the track element upon which you double clicked will open. Double clicking on the sonogram will open the Bpf corresponding to the last selected treatment.

The "Replicate in Time. . . " (Edit menu) and "Invert" (Treatment menu) are activated. The Invert function will carry out transposition in the opposite direction.

19.4.3 "Dynamic TimeStretching"

In the "Treatments" menu, select "Add Dynamic TimeStretch. . . " this is done the same way as for Dynamic Transposition.

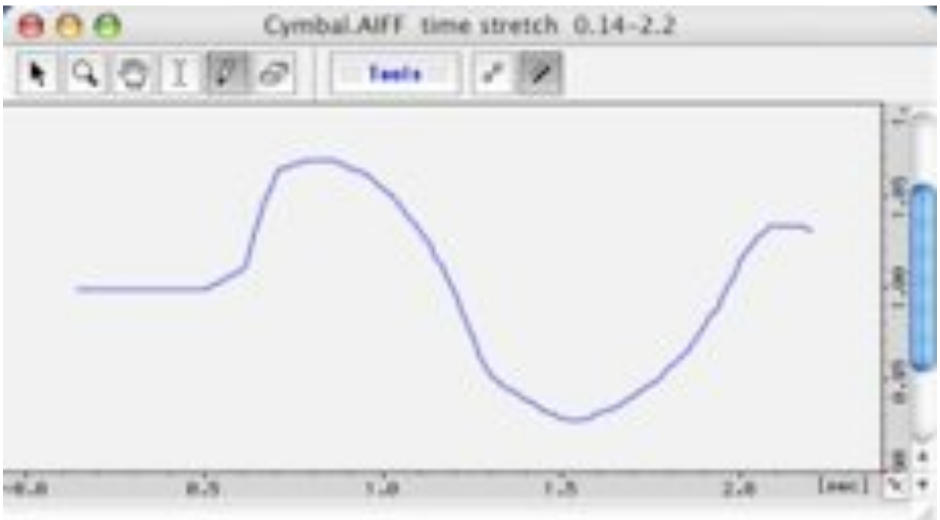

The horizontal axis is time and the vertical axis is the stretching/compression coefficient. The editor limits the compression coefficient to 0,01 and the stretching coefficient to 10.

The corresponding track element will be positioned on the sequencer track.

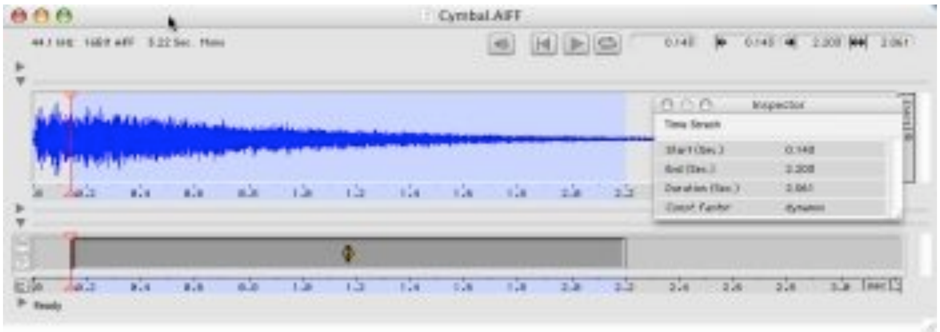

To edit the Bpf, double click on the track element (or on the corresponding sonogram surface). Modifications are carried out as for all Bpf's:

The Inspect toolbox allows access to the following parameters: start, end and duration.

The "Replicate in Time..." (Edit menu) and "Invert" (Treatment menu) are activated. The Invert function will generate an inverted Bpf.

19.4.4 "Breakpoint for Gain"

You can modify signal gain using the "Add Breakpoint for Gain..." item in the Treatments menu. The horizontal axis is time and the vertical axis is the gain in dB (logarithmic amplitude). The editor limits the gain to +116 dB and the attenuation to -116 dB.

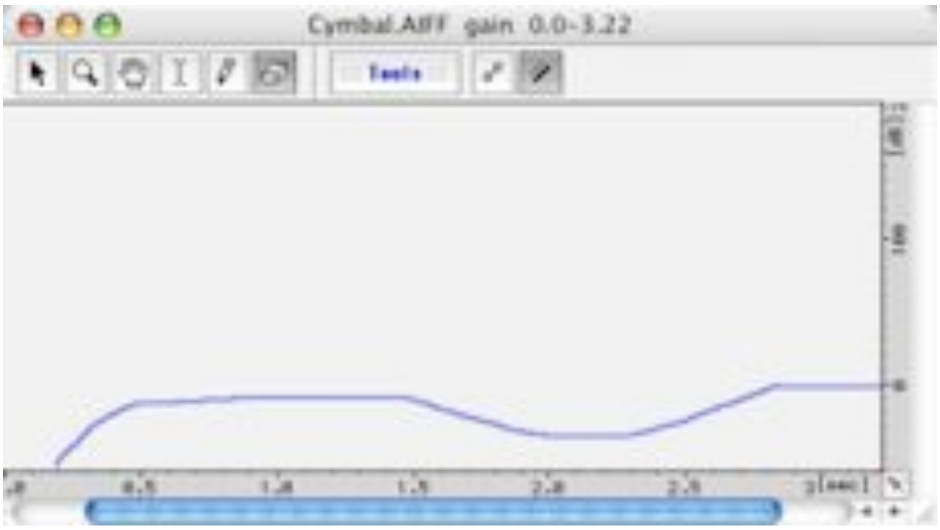

Note: It is advisable to check that normalization is activated when augmenting the gain.

The corresponding track element will be positioned on the sequencer track.

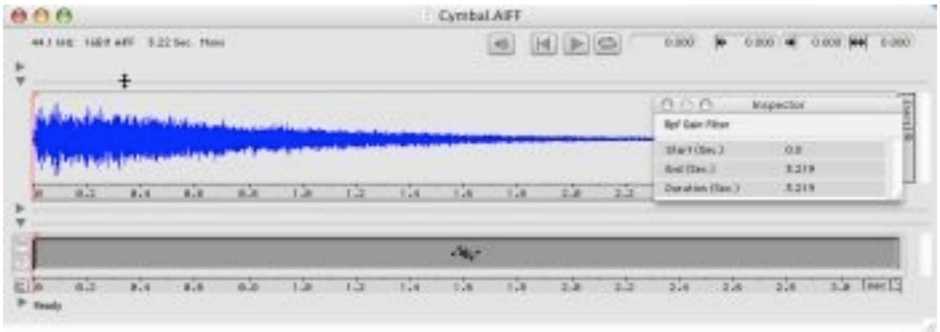

The Inspector toolbox allows access to the following parameters: start, end and duration.

To edit the Bpf, double click on the track element (or on the corresponding sonogram surface). Modifications are carried out as for all Bpf's:

The "Replicate in Time. . . " (Edit menu) and "Invert" (Treatment menu) are activated. The Invert function will generate an inverted Bpf.

19.4.5 "Breakpoint Filter"

In the "Treatments" menu, select "Add Breakpoint. . . " to open the Bpf editor which allows you to draw and manipulate the Bpf that represents your filter.

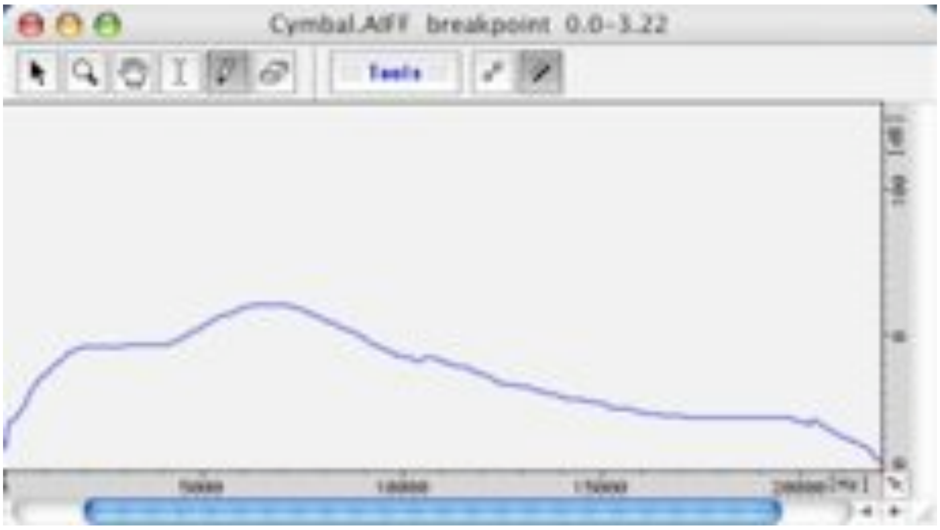

Frequencies are shown along the horizontal axis and the amplitude (logarithmic amplitude in dB) along the vertical axis. The editor limits the gain to +116 dB and the attenuation to -116 dB.

> **Note:** If the Bpf resolution is too low, especially for Bpf filtering (Breakpoint Filter), simply change the sample rate using the Resampling item in the Bpf menu.

At the same time a rectangle with the same duration will appear on the sequencer track, carrying the icon that corresponds to the treatment (track element).

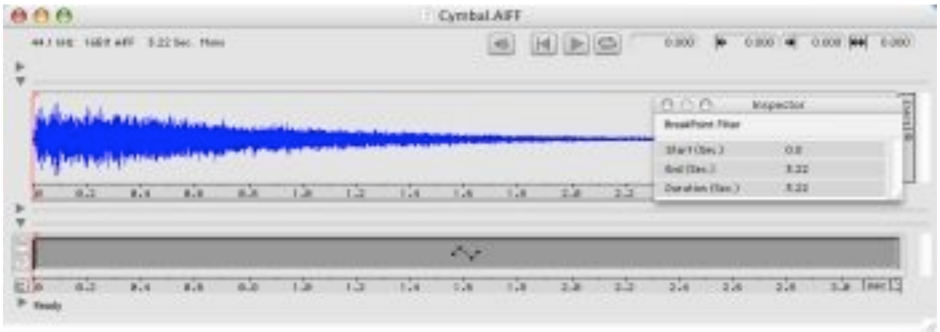

To edit the Bpf, double click on the track element (or on the corresponding sonogram surface). Modifications are carried out as for all Bpf's:

You can move or resize along the time axis, or you can duplicate the rectangle (and thus duplicate the treatment) on the track in the same way as with the other elements (see section 19.5).

The Inspector toolbox allows access to the following parameters: start, end and duration.

The "Replicate in Time..." (Edit menu) and "Invert" (Treatment menu) are activated. The Invert function changes the gain sign (opposite Bpf).

19.4.6 "Dynamic Formant"

This treatment is not yet operational.

19.5 What all the treatments have in common

19.5.1 Normalisation

Normalization at treatment output is not active by default. Therefore, exercise care when applying any treatment that increases gain.

Activate normalization in "Processing Parameters" via "Process Treatments..." or "Process Selection..." in the Processing menu (see section 25.1.2).

Normalization is active by default for "Source Filter Synthesis" and "Generalized Cross-Synthesis" (see sections 23.1 and 23.2). You can deactivate it in the corresponding control panels.

> **Note:** Real Time Processing behaves differently. Please consult section 28.

19.5.2 Important note

The following does not apply to special treatments (see section 23). The special treatments are: Source Filter Synthesis, Generalized Cross-Synthesis, Normalize and Samplerate Conversion in the Processing menu.

19.5.3 Reminder

Moving, duplication, deleting and saving via drag and drop can be applied [simultaneously] to several treatments of any type.

To process one or two treatments of any type, you must use the sequencer's track elements.

Some parameters can be edited directly via the Inspector toolbox.

19.5.4 Selecting several treatments

In order to select several treatments, simply click on the track elements with any tool (except the marker tool) while holding down the **Shift** key.

In the Edit menu, the Select All Treatments item allows you to select all defined treatments.

19.5.5 Deleting several treatments

To delete one or several selected treatments, simply press the **Delete** or **Backspace** key, or select Clear in the Edit menu.

19.5.6 Moving treatments

You can move the selected treatment(s) horizontally via one of the track elements selected with any tool (except the marker tool).

19.5.7 Duplicating treatments

Press the **Option** (**Alt**) key while using the pointer (or grab) tool to duplicate the selected treatment(s), and to place the duplicate at the desired point on the track.

Note: Image Filter will appear to be duplicated, but in fact is not (the image will not display).

19.5.8 Resizing treatments

To resize a treatment, place the tip of any tool (except the marker) on a track element. It will turn into 2 small black triangles. Then click and horizontally enlarge or reduce the treatment, depending on the edge you click on.

Note: This action can only be applied to one treatment at a time.

19.5.9 "Replicate in Time. . . "

This function, in the Edit menu, allows you to duplicate selected treatment(s) along the time axis.

Choose the number of looped copies required, and their time shift to the right (positive time in seconds) or to the left (negative time in seconds).

Of course, you can only check the Align to Markers box if there are markers. It allows you to align the copies with the markers at a spot later (or equal) to the desired time shift.

> **Important:** If the number of copies requested is greater than the number of usable markers, the additional copies will be superimposed and aligned with the last marker.

19.5.10 "Expand To Markers"

This function is in the Edit menu (keyboard shortcut **X**) and allows you to extend the selected treatment(s) to the nearest left and right markers.

19.5.11 "Expand To Grid"

This function is in the Edit menu (keyboard shortcut **Shift+X**) and allows you to extend the selected treatment(s) to the nearest left and right grid references.

19.5.12 "Invert"

This function is in the Treatments menu, and will only be active if appropriate to the treatment in question.

Thus, it does not apply to:

- "Clipping"
- formantic filtering ("Constant Formant")
- "Freeze"
- "Reverse/Repeat"

According to each case, it acts upon the gain or upon another parameter.

- When used upon surfaces: it changes the gain sign.
- When used upon transpositions: it changes the transposition sign, or generates the opposite Bpf.
- When used upon "TimeStretching": it gives the inverted coefficient (y = 1/x) or the opposite Bpf.
- When used upon Bpf filtering ("Breakpoint Filter"): it changes the gain sign.
- When used upon Image filtering ("Image Filter"): it generates the negative version of the image.
- When used upon Band filtering ("Band Filter"): it toggles between Band Pass and Band Reject.
- When used upon "Breakpoint for Gain": it generates the opposite Bpf.

19.5.13 "Drag and Drop"

Drag and Drop several selected treatments on the desktop to save all their characteristics (horizontal and vertical positions, gain, coefficients, Bpf's etc): the entire set will be named "AudioSculpt extract" or the "AudioSculpt.txtClipping".

To re-use this treatment or set of treatments in an identical way upon any sound, simply drag and drop the AudioSculpt.textClipping extract to the sound's sonogram (see section 5.2)

You can give any desired name to the extracts.

> **Note:** The extracts are automatically and incrementally numbered.

19.5.14 Multiple treatments

Note that treatments can overlap. They can be covered over and defined in any order. AudioSculpt will use them to process the sound in a single pass.

> **Note:** It is preferable to carry out the fewest number possible of processing passes in SuperVp. Therefore, it is a good idea save your treatments (see section 19.5.15) in a single set, in order to edit them later on if required, and apply them to the sound again (or to another sound).

> **Note:** Think carefully before overlapping treatments whose effects may be partly or completely contradictory.

19.5.15 Saving files and using saved files

A treatment file contains all the treatment elements defined for use upon a sound (the entire set of track elements). When you close an AudioSculpt window, or quit the application, a dialog box will ask if you wish to save the treatment, and choose a name (sound_name.trt), as well as the location to which it must be saved. The default folder is called "Treatments", and is in the application folder. You can of course choose any other folder. It is advisable to keep the ".trt" extension.

This allows you to load the entire set in the sound's AudioSculpt window when re-opening. It also allows you to apply the treatment set to any other sound. To do this, after loading a sound, go to the File menu, then Open Treatment…, then choose a ".trt" file in the standard dialog box. Treatments will be displayed in the window, and where applicable, be added to existing treatments there.

You can save this file type at any point in the session: Go to the File menu, then Save Treatments or Save Treatments As… and choose the name and location (default: "Treatments" folder) of the new file.

19.6 Common features in transpositions

19.6.1 Preserving transients

See section 27 : "Detecting and preserving transients".

19.6.2 Time correction

Each transposition (no matter how many) that has been defined for a sound can be carried out with or without time correction. This also applies to time variable transpositions (Dynamic Transposition). Use the Inspector window to choose this parameter for each treatment (toggle on/off opposite "Time Correction).

For constant transposition, you can also make your choice in the usual dialog box.

19.6.3 Spectral envelope

You can carry out time corrected transpositions while preserving the spectral envelope. Use the Inspector window to choose this parameter for each treatment (toggle on/off opposite "Time Correction). This option is checked by default.

20 Multichannel sounds

20.1 Stereophonic or two-channel sounds

AudioSculpt works in the same way as for mono sounds.

The upper zone displays the sum of the channels. The sound is played in stereo. AudioSculpt displays both channels, and can generate sonograms for them. It can also generate spectrograms for them. The tuning fork shows both spectrums (in blue). Reminder: the spectrum is red at the point where the cursor is located.

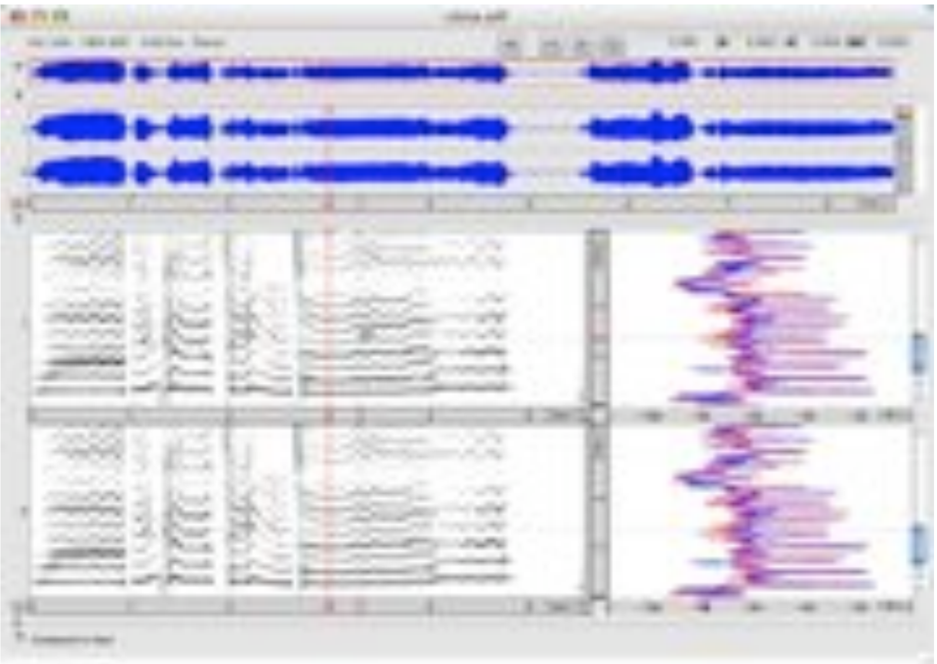

For zones 2 (window into the sound) and 3 (sonogram) you can hide either of the channels using the button at the left of the rulers.

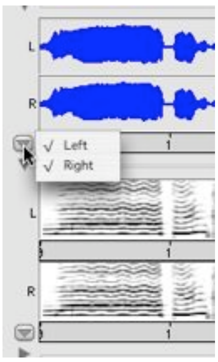

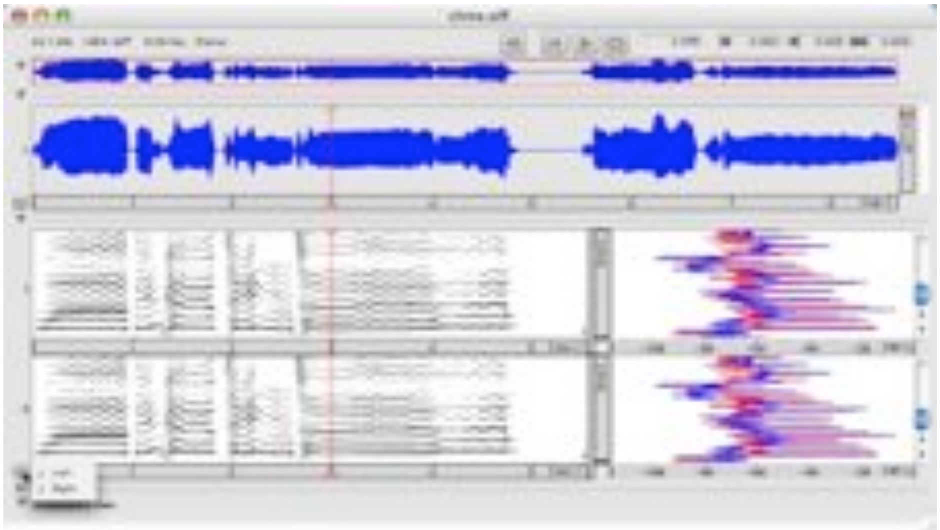

Treatments can be applied to either or both of the channels. This can be set in the following dialog boxes: Process Treatment. . . , Process Selection. . . , and Real-time Processing Settings. . . of the Processing menu.

Some analysis types apply to both channels or only to one of them (left or right). There are others that can only be applied to one channel at a time. The settings are in the analysis dialog boxes.

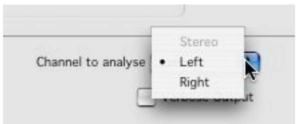

Stretching one channel will add silence to the end of the other. Compressing one channel will add silence to its beginning. Exactly the same applies to a selected part.

Freezing one channel will add silence to the end of the other.

20.2 Multichannel sounds (other than stereo or two-channel)

They will be processed in the same way as stereo sounds. **HOWEVER**, AudioSculpt can only play the first two channels on the Mac itself. But if there is a sound card that allows it, all the channels will be played.

For zones 2 (window into the sound) and 3 (sonogram) you can hide either of the channels using the button at the left of the rulers.

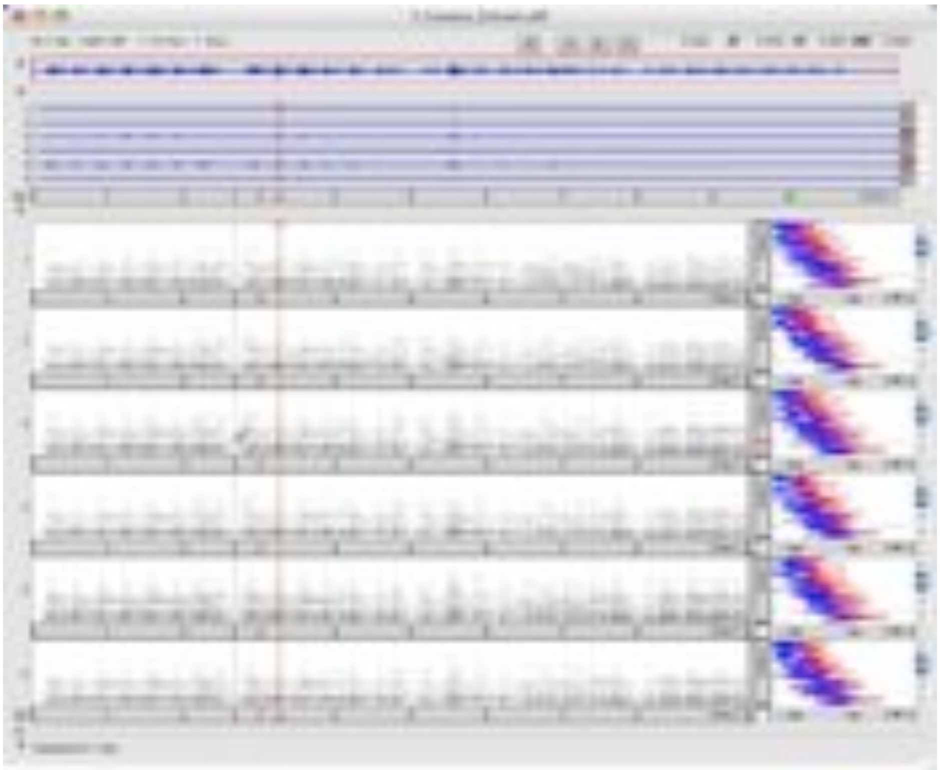

Treatments apply to all channels or to one channel, as desired. Certain analysis types apply to all channels or to one channel, as desired. There are others that can only be applied to one channel at a time. The settings are in the analysis dialog boxes.

Stretching one channel will add silence to the end of the other. Compressing one channel will add silence to its beginning. Exactly the same applies to a selected part.

Freezing one channel will add silence to the end of the other.

21 The Sequencer

21.1 The sequencer zone

In the AudioSculpt window, the sequencer zone (zone 5) is located under the sonogram zone.

The little triangle on the left shows that this zone is unfolded. When the AudioSculpt window opens, it [initially] only displays one (empty) track.

The sequencer tracks display the treatments, including duration and time layout.

21.2 The sequencer tracks

When you define a treatment (except "Source Filter Synthesis", "Generalized Cross-Synthesis", "Normalize" and "Samplerate Conversion" in the Processing menu), a rectangle (track element) with the same duration will be displayed on the track at the desired time point.

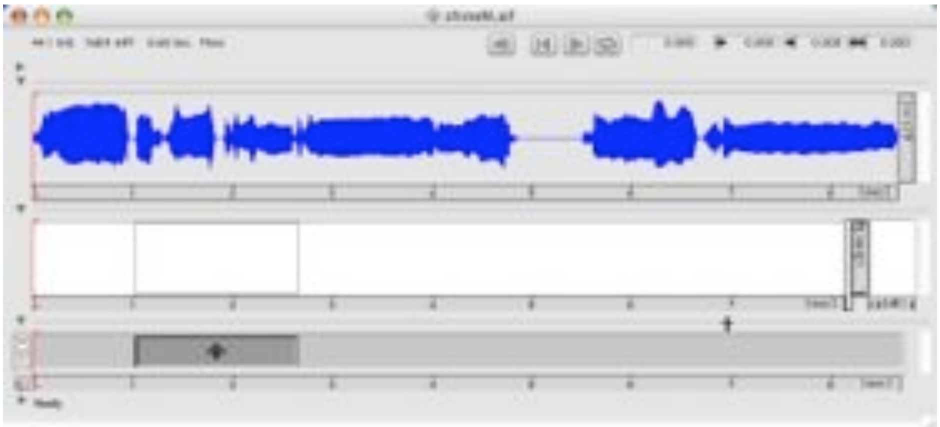

If you successively define several treatments, they are added to the same track.

You can create as many tracks as you like using New Track in the File menu.

You can force track elements onto a given track: Simply select the desired track and click on it with any tool (except the marker). It will turn dark gray.

If "Create Tracks When Needed" (Treatments menu) is checked, a new track is only automatically created if necessary (i.e. so that one treatment will not hide another). This is especially useful if you use many filtering treatments (especially with the pencil tool).

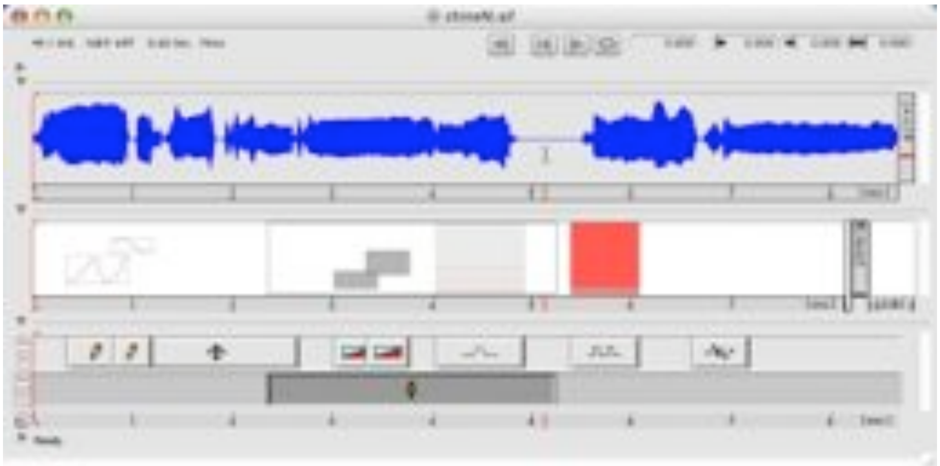

You can delete a selected track (and its contents) using Delete Selected Track in the Edit menu.

It is also possible to do the same action using the **ctrl** key. In the contextual menu choose "Delete This Track"

Note: You cannot delete all the tracks. At least one track remains.

You can move any track element to [on] any track, using any tool except the marker. This can be done with any number of selected elements.

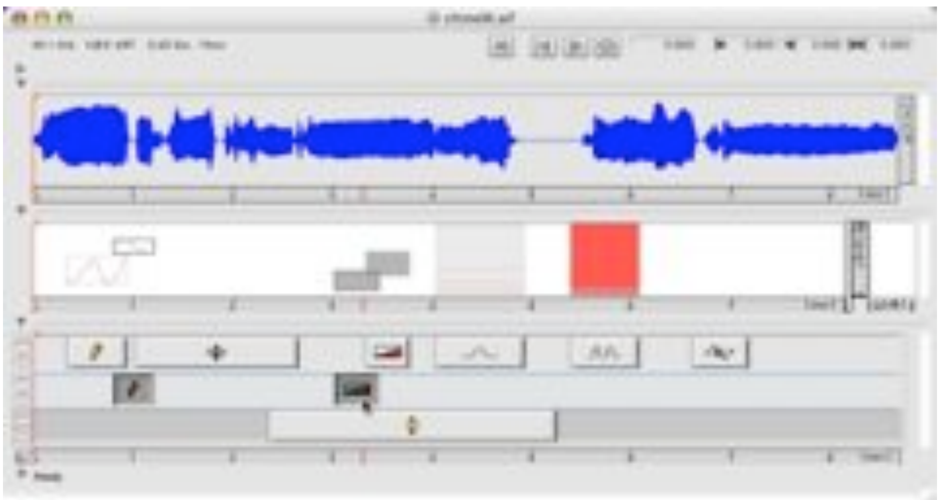

To preserve the time position of an element while moving it to another track, simply hold down the **Shift** key. This can be done with any number of selected elements.

This is the track heading display:

Each track has 2 buttons on the left: the mute and solo buttons for each track. They are the same as those found in any multi-tracking software.

- "mute" : deactivates the track (and therefore the treatments on that track).
- "solo" : activates the track and deactivates all other tracks.

> **Note:** The mute function overrides the solo function.

The button (a small magnet) is bottom left (opposite the track rulers). It is shared by all track elements and/or treatments on the sonogram. It allows you to activate or deactivate (default) the snap ("magnetism") function , and to set the desired type for all the tracks.

Time and/or frequency can be set, with or without offset. Move one or several track elements and /or treatments on the sonogram, and it/they will automatically align with the nearest marker, according to options you chose previously.

Along the horizontal axis (time), elements can "snap" to:

- The markers: treatment start or end will snap to the nearest marker.
- To the grid references (see section 22): treatment start or end will snap to the nearest

marker.

- To the treatments ("Glue To"); the treatment will snap to the nearest treatment.

Along the vertical axis (frequency), the element will snap to the nearest grid frequency reference (see section 22) above or below, or will center between them. Snapping will only be applied to surfaces being processed (moved) on the sonogram.

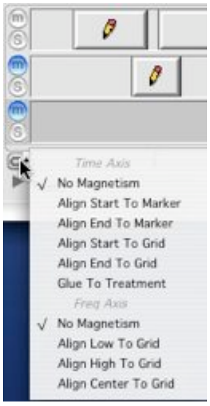

Note: Snapping will override time position preservation when moving en element (treatment) from one track to another.

Expand To Markers in the Edit menu (keyboard shortcut **X**) and Expand To Grid (keyboard shortcut **Shift+X**) expands the selected treatments to the nearest outer markers or time grid graduations respectively.

Track information is saved in the treatment file.

22 The Grid

You can display a reference grid on the sequencer and on the sonogram: Choose Show Grid Settings... in the Window menu, and the following toolbox will be displayed:

The **tab** key allows you to hide/mask this toolbox as well as three others: Tools, Sonogram Display, and Inspector by checking them.

Check the Time Grid box to display time references, which you can parameterize: Seconds, Bpm, Or Samples. You can also include an offset in Seconds, Bpm, Or Samples. The offset allows you to time shift the grid away from zero time (e.g. you can shift it to an attack).

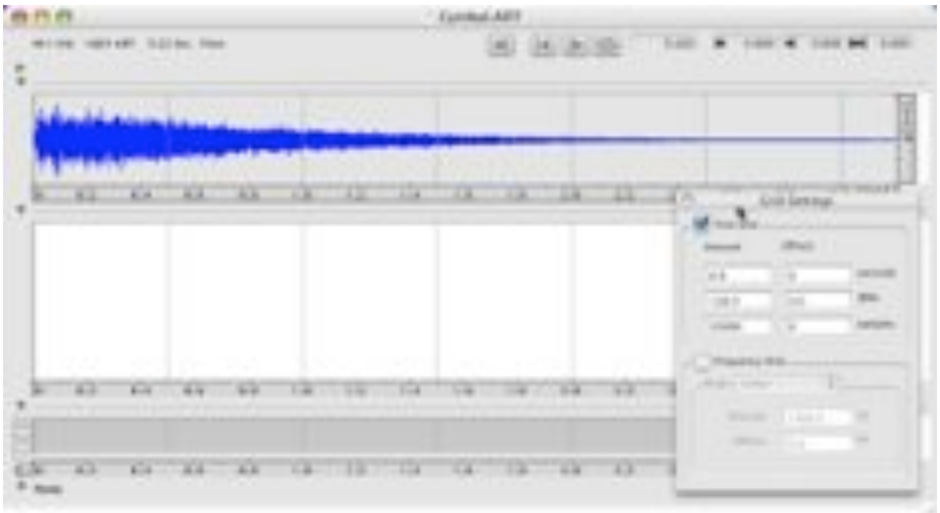

Check the Frequency Grid box to display frequency references. You can parameterize them. First choose the mode (linear or logarithmic), then the resolution in Hz, and then the Offset value, in Hz.

- In linear mode, resolution and Offset are in Hz,
- In logarithmic mode (Harmonic mode), resolution is in cents or in intervals, and the Offset becomes the Base Frequency in Hz.

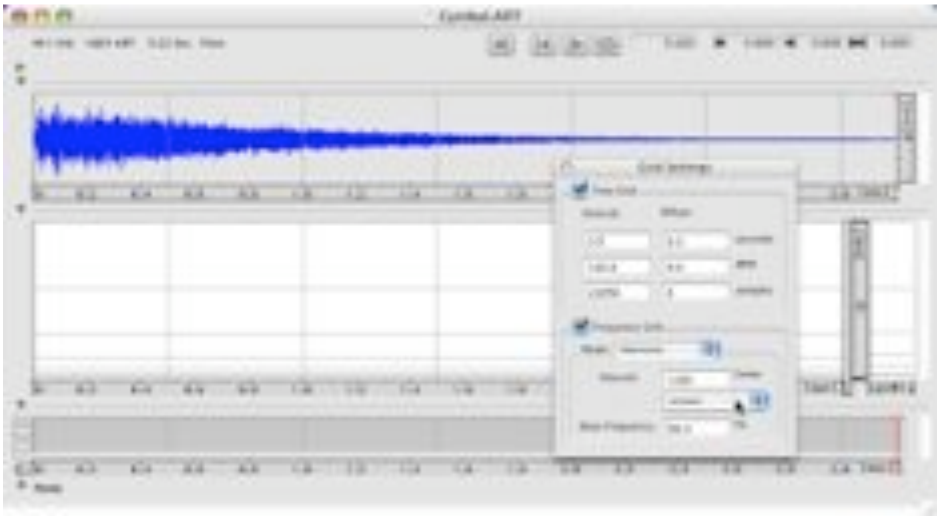

The snap to grid function is activated in the sequencer (see section 21.2).

23 Other treatments

The following treatments are accessed via the Processing menu:

Note: These treatments are not written onto the sequencer tracks, and cannot be saved. Therefore, processing is carried out individually and directly for each of them.

23.1 Source Filter Synthesis :crossed source/filter synthesis

Typically, crossed source/filter synthesis is carried out by multiplying the FFT spectrum of a sound with another sound’s LPC analysis generated spectral envelope. Thus, the first sound is filtered using the second sound’s spectral envelope.

As regards the filtering sound (the second sound) you may choose from among several analysis types:

- Default: LPC.
- FFT (which is in fact, general crossed synthesis, see below).
- Discrete Cepstrum.
- Inverse LPC.
- Inverse Discrete Cepstrum.

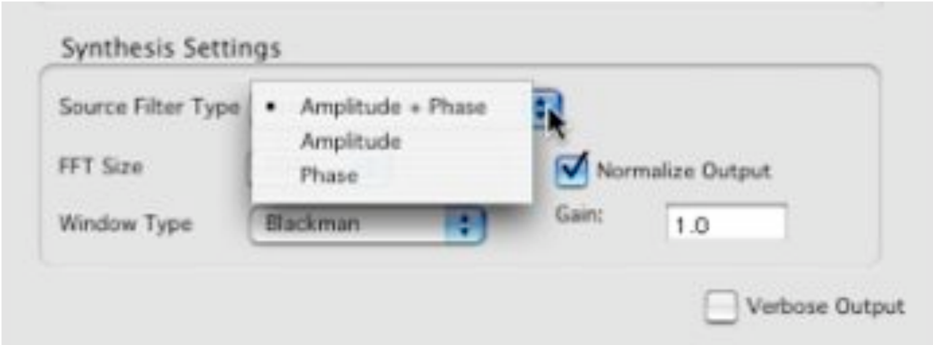

After this, you have the following Source Filter Type options:

- Amplitude+Phase
- Amplitude
- Phase

Choose the Source Filter Synthesis option found in the "Source Filter Synthesis. . . " item, via the Processing menu.

> **Note:** Discrete Cepstrum and Inverse Discrete Cepstrum do not process phase. Only amplitude is processed.

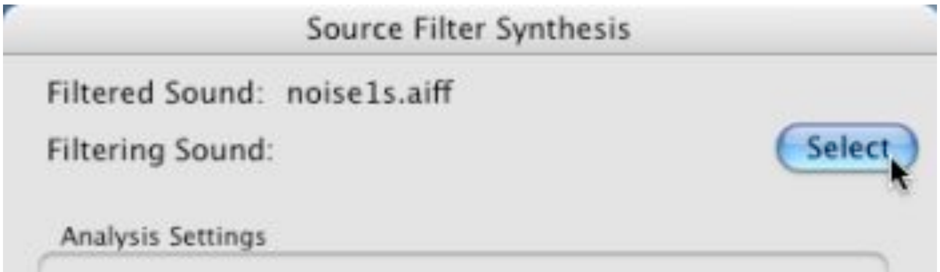

The opened (or foreground) sound is filtered.

This toolbox displays the filtered sound's name, and invites you to choose the Filtering Sound by clicking on the Select button.

A standard dialog box opens, allowing you to navigate through the disk drive tree and choose a filtering sound.

When you have selected the filtering sound, set the usual analysis parameters for each of the 2 channels: (Analysis Settings): Window Size, Window Step, Window Type as well as the number of LPC or discrete Cepstrum analysis poles (Analysis Order). The more poles there are, the more detailed the spectral envelope will be. You might want to use only two poles, in order to model only the most important forms of the spectral envelope.

The Synthesis Settings: Source Filter Type, FFT Size and Window Type. For more details, consult section 25.1. The adjustable gain is intended for real-time use.

Note: The output is normalized by default.

Verbose Output is unchecked by default. If you check this option (for advanced users: flag-v), all the information generated by SuperVp, apart from the command line, will be displayed in the SuperVp Console, if it is unfolded (see section 24.1).

After setting your parameters, press **Enter**.

The settings panel will close and a standard dialog box will open, inviting you to choose a name and a location for the result (default: the Sounds folder). After choosing, press OK. The window will close, and processing begins.

When processing is done, a new AudioSculpt window will open showing the result (this window will have the same name as the new sound).

Note: The [display of the] filtering sound takes on the same size as the filtered sound.

Here is the result of "noise1.aiff" filtered by the sound Cymbal.AIFF with LPC analysis:

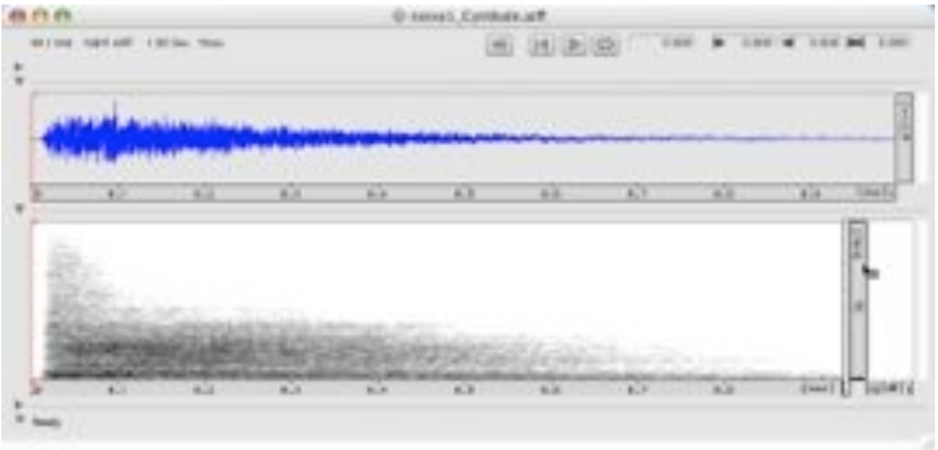

The sound "noise.aiff" filtered by "Cymbal.AIFF" with inverted LPC analysis:

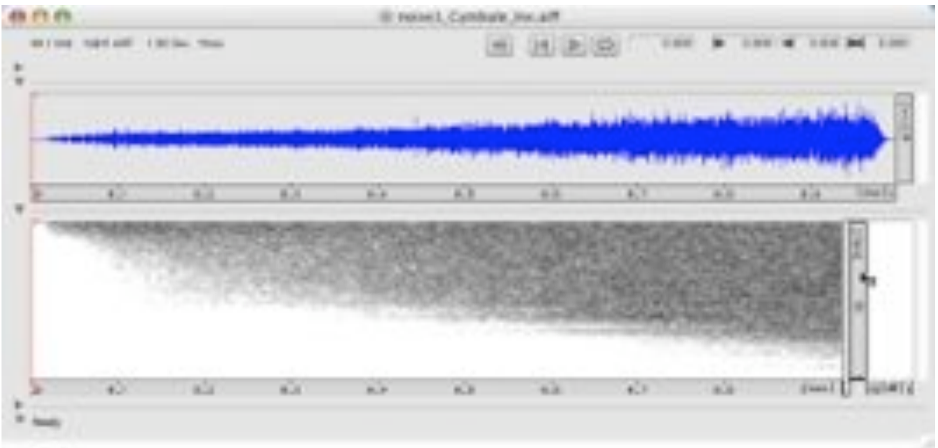

23.2 Generalized Cross Synthesis

This carries out FFT analysis of both sounds (i.e. amplitude and phase).

The opened (or foreground) sound is sound 1.

In Processing, choose Generalized Cross Synthesis. The next palette (or toolbox) will open, reminding you of sound 1's name (Sound 1), and will invite you to choose Sound 2, by clicking on the Select button.

A standard dialog box opens, allowing you to navigate through the disk drive tree and choose a sound.

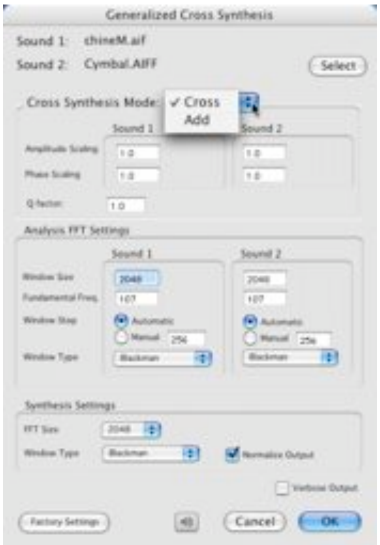

You can choose between 2 modes:

- "Cross" : cross synthesis (analogous to the previous process); you need to set the amplitude coefficients (Amplitude Scaling) and Phase Scaling pour each of the sounds as well as the coefficient of the product of the "Q Factor" amplitudes.
- "Add": Addition of the sounds (mixing); here you only need to set Sound 1's Amplitude Scaling (it allows you to vary the amplitude proportions).

Set the usual FFT and Resynthesis parameters. Please refer to section 25.1.

> **Note:** The output is normalized by default.

"Verbose Output" is unchecked by default. If you check this option (for advanced users: flag-v), all the information generated by SuperVp, apart from the command line, will be displayed in the SuperVp Console, if it is unfolded (see section 24.1).

After setting your parameters, press **Enter**.

The settings panel will close and a standard dialog box will open, inviting you to choose a name and a location for the result (default: the Sounds folder). After choosing, press **OK**. The window will close, and processing begins.

When processing is done, a new AudioSculpt window will open showing the result (this window will have the same name as the new sound).

> **Note:** The resulting sound will be of the greater of the lengths.

23.3 Normalization

Normalization is applied to the entire sound. The Normalize… function is in the Processing menu. It will open a dialog box allowing you to set the margin you desire Normalization to leave (in dB).

After setting your parameters, press Enter. The settings panel will close and a standard dialog box will open, inviting you to choose a name and a location for the result (default: the Sounds folder). After choosing, press **OK**. The window will close, and processing begins. When processing is done, a new AudioSculpt window will open showing the result (this window will have the same name as the new sound).

Normalization will be carried out In-Place if you check the appropriate box.

> **Important:** In this case, it is applied to the original sound, and therefore cannot be cancelled.

23.4 Samplerate Conversion

Samplerate Conversion... in the Processing menu will open a dialog box allowing you to convert the Samplerate, as well as to change the Quantize resolution (Samplesize), and the file type.

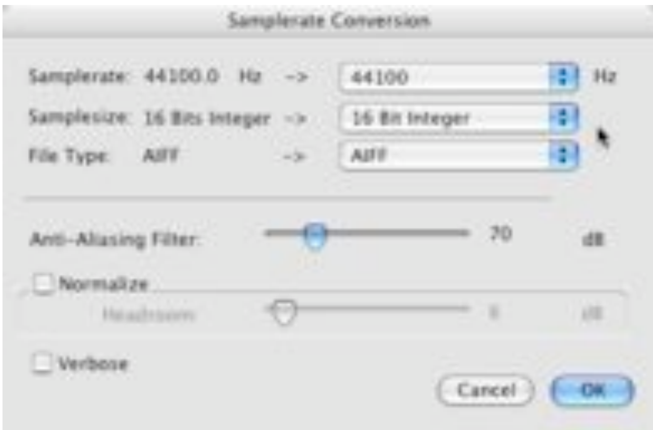

The Anti-Aliasing filter is adjustable. It avoid the distortion that results from folding.

It is a high-order lowpass filter. The greater the high-frequency attenuation, the lower the distortion. However, this trades off against more processing cycles. The default value (70

dB) is the folded frequencies' minimum attenuation value.

This conversion does not go via the usual settings panel. Instead, it offers:

- adjustable normalization
- the verbose option. If you check this option (for advanced users: flag-v), all the information generated by SuperVp, apart from the command line, will be displayed in the SuperVp Console.

After setting your parameters, press **Enter**. The settings panel will close and a standard dialog box will open, inviting you to choose a name and a location for the result (default: the Sounds folder). After choosing, press **OK**. The window will close, and processing begins.

When processing is done, a new AudioSculpt window will open showing the result (this window will have the same name as the new sound).

24 The Consoles

There are two types of console:

24.1 The integrated console

The console attached to the AudioSculpt window (zone 5, at the bottom of the window) which belongs to the opened sound file. Please refer to section 5.2.

It is folded in the beginning, and displays "Ready": Thus is the SuperVp console, in which the command lines are written, as well as any alerts generated by SuperVp or pm2. To unfold it, click on the triangle above it.

If the Verbose Output or Verbose is checked (for advanced users: flag-v, in the command panels, the console will display detailed information from SuperVp or pm2.

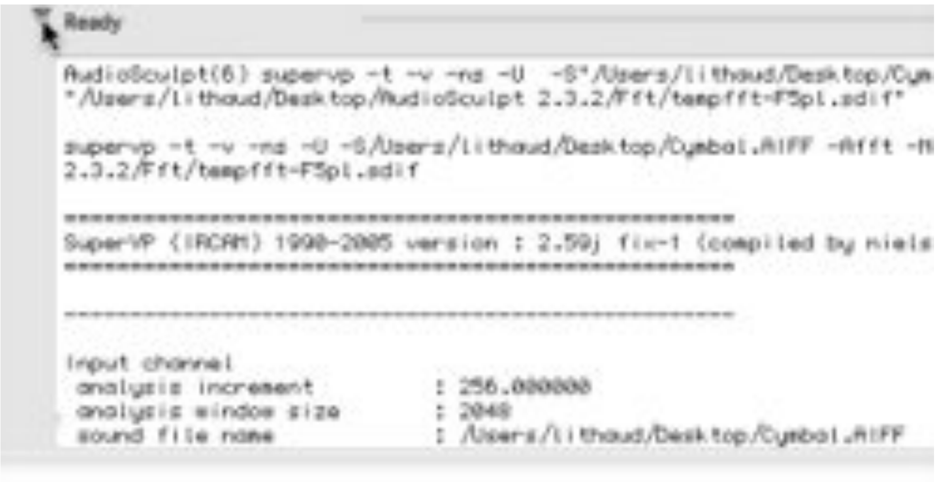

This console can be saved via Save Console As. . . in the File menu.

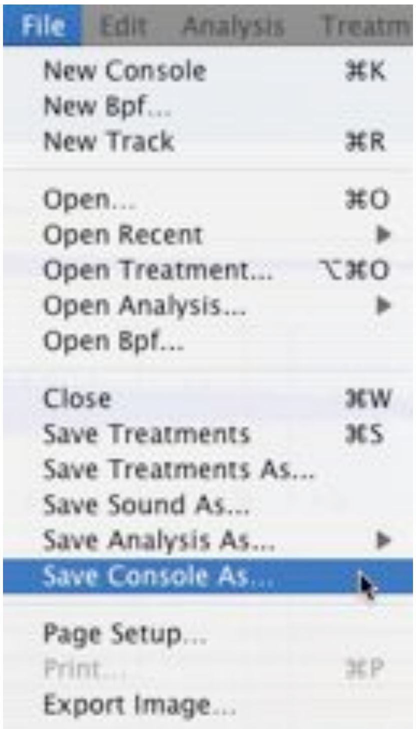

24.2 The empty console: New Console

This is obtained via New Console in the File menu. It allows you to launch command lines. You can saving this console (when active) via File menu, then Save Console or Save Console as. . . These menu items only appear in the menu in the following case: console active.

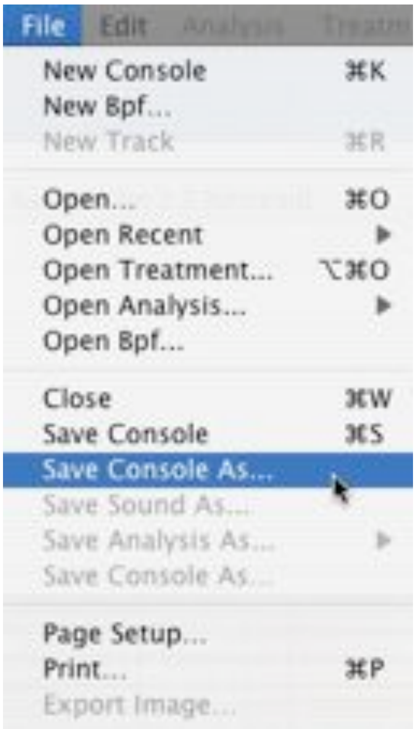

These consoles are saved in text format and can be opened in any text editor.

Print in the File menu allows you to print any active console.

25 Processing

25.1 Processing the entire sound "Process Treatments. . . "

In order to start processing, choose Process Treatments. . . in the Processing menu. This command remains grayed out if there is not at least one treatment defined. The processing takes into account all the treatments defined in the sequencer, except those that are on tracks rendered inactive by the mute button (see section 21.2).

Tip: Reminder: The mute function overrides the solo function.

A dialog box opens and allows you to set the desired parameters:

25.1.1 Analysis parameters: FFT Settings

- "**Window Size**" : Window size corresponds to the number of samples in each sound analysis. Window size is the essential analysis parameter: it determines the frequency and time resolutions. The **Window Size** and **Fundamental Frequency** fields are interdependent.
- "**Fundamental Frequency**" in Hz : allows you to directly adjust the frequency resolution. For a mono sound, this is its fundamental frequency. Window size, step size and FFT size will be adjusted accordingly. Window size is programmed to be 5 times the length of the period corresponding to the fundamental frequency shown.
- "**Window Step**" - Window Step: this sets the time interval, measured in samples, be-

tween 2 successive analyses. In automatic mode, SuperVp adjusts the step size to give an optimal result (it can make the step size variable if necessary) according to the defined treatments. This is the recommended mode. In manual mode, you set the step size yourself (for advanced users: flag **-I**).

> **Important:** automatic mode works differently for analysis (see section 16.1).

- "**FFT Size**": sets the number of analysis points applied to the samples. It is necessarily equal or greater than the window size.
- "**Analysis Window**" : Allows you to set the window type used in the analysis. Three window types are offered: **Blackman**, **Hanning** and **Hamming**.

25.1.2 Synthesis Settings

- **Phase Synchronous Processing**: carries out phase synchronization (for advanced users: flag -P) ; this process usually improves the result. Do not uncheck Phase Synchronous Processing unless you know what you are doing.
- **Preserve Transients**: only active if the Phase Synchronous Processing box is checked (see above).
- **Analysis Window** : Allows you to set the window type used in the analysis. Three window types are offered: Blackman, Hanning and Hamming.
- **Filter Superposition Mode** : offers the choice between Multiply and Maximum modes for overlapping treatments that alter the gain (surface filters and formant filters).
  - mode "Multiply" : dB gain values are accumulated.
  - mode "Maximum" : the greater gain value is taken as the maximum gain.
- **Envelope Preservation Order** : sets the number of poles is spectral envelope preservation is requested for a transposition.
- **Normalize Output** : Normalization at treatment output is not active by default. Therefore, exercise care when applying any treatment that increases gain.
- **Transient Preservation Settings** : allows you to make settings necessary for transient preservation (see section 27).

Verbose Output is unchecked by default. If you check this option (for advanced users: flag-v), all the information generated by SuperVp, apart from the command line, will be displayed in the SuperVp Console. They will be visible if it is unfolded (see section 24.1).

Click on Process. AudioSculpt will ask you for the name (here the default name is africa.aiff.out) and the location where the file must be saved (default is the Sounds folder in the application folder). Press OK and processing begins. The progress bar at the bottom of the AudioSculpt window, allows you to follow processing progress. When this is done, a new AudioSculpt window will open, displaying the processed sound.

> **Note:** If you carry out several processes using the same name, AudioSceltpt will suggest an automatic incremental numbering of the new sound name.

25.2 Process Selection. . .

If desired, you may process only part of the sound with the corresponding treatments. After selecting the desired part, choose Process Selection. . . in the Processing menu. This command remains grayed out if there is not at least one part of the sound selected.

25.3 Temporary files

Files generated by AudioSculpt for SuperVp are temporarily stored in the Temp folder. They are immediately replaced when you carry out a new processing instance of the same type, and are irretrievably deleted when you quit AudioSculpt. (It is not necessary to preserve them, they will be generated again if you carry out the same process afresh, and if you have saved the treatment).

However, advanced users will need to move them to another folder if they desire to examine them (they are text files, and can be very long, edited for example in BBEdit), or if they wish to edit them for later use in a command line.

25.4 Work session

Options that are chosen for Processing remain valid for the entire session (until you quit AudioSculpt). However, they can be saved in the Preferences using the Save button there. The Reset Parameters button allows you to come back to these values after any change. The Factory Settings button resets the software's factory default values.

26 Interrupting Processing

You can always stop processing by clicking the Stop button on the console. The processed part of the sound will be saved, and the new sound's window will open, even if normalization is active. The same applies to analysis.

27 Detecting and preserving transients.

In the setting panel that is opened by Process Treatments. . . , Process Selection. . . or Real-time Processing Settings. . . (Processing menu), the Preserve Transients option provides transient detection, for the purpose of improving certain treatments. This option uses the parameters that were set in the Transient Preservation Settings part. Transient detection can only be carried out in Phase Synchronous Processing mode (see section 25.1.2).

Naturally, these parameters are also to be found in the transient detection marker settings panel (Segmentation Type: Transient Detection), in the Analysis menu (Generate Markers. . .).

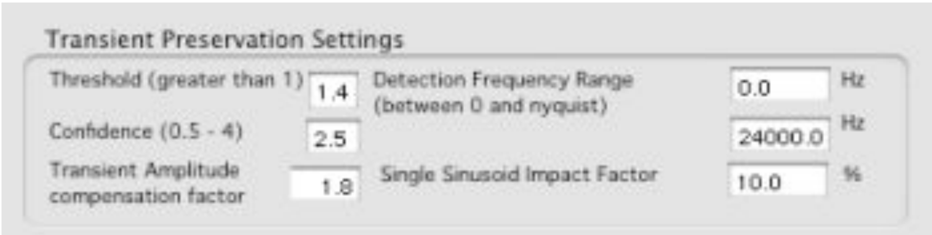

Transient detection is based on the energy concentration in the right hand part of the analysis window. In order to effect detection in polyphonic signals, energy is measured in the frequency bands.

During processing, the analysis window moves along the sound. When the start of a sound event is placed in the right hand part of the window, this is considered to be a transient situation. In borderline cases, the start of a sound event is placed in the middle of the window. Assessment of energy concentration is carried out on a normalized basis, so that values will lie between 1 and 10.

In an ideal situation, i.e. when there is no noise or residual sound, a transient will generate an energy concentration curve that starts at 10 (i.e. at the extreme right of the window) and decreases down to 1 (borderline case: middle of the window).

In the real world however, the presence of noise reduces the normalized located energy maximum value. The higher the amplitude of a sound event's start in proportion to the noise present, the closer will be the normalized located maximum energy value to 10. It is not only transients that cause energy displacement, but also signals with noise content.

In order to distinguish between the two different types of cause, we use a statistical model.

- **Threshold (larger than 1)(1-10)** : the located normalized energy threshold beyond which transients may be detected. The default value of this threshold is 1.4. For a value of 10, nothing is detected. For a value of 1, it is very likely that transients will be detected, even though they are nothing more than variations in residual noise. This threshold may be lowered in order to increase detection sensitivity and vice versa.
- **Detection Frequency Range (between 0 and nyquist)** in Hz : sets the frequency limits (minimum: 0Hz and maximum: Nyquist, sampling frequency divided by 2) for bands used in located energy evaluation as well as for detection. You can modify this value, but for processing, the entire signal band is used.
- **Confidence (0.5-4)** : confidence degree threshold, deduced from the statistical model, required for a transient. The default value is 2.5. This threshold may be lowered in order

to increase detection sensitivity and vice versa.

- **Single Sinusoid Impact Factor** : coefficient in percent that sets the frequency band size in proportion to a stationary sinusoidal. The energy location is evaluated separately for each band. We recommend strongly that you do not modify this value, unless you really know what you are doing. The default value is 10%.
- **Transient Amplitude compensation factor** : the value of 1.4 compensates the amplitude loss caused by the transient preservation algorithm. By greatly increasing this value (e.g. 2000) and above all, not omitting to normalize, only the transients will remain.

28 Realtime mode

AudioSculpt (SuperVp in fact) can play all or part of its treatments in real-time. Naturally, this depends on computer power, and the number and type of the desired treatments. The mute and solo functions allow you to deactivate some of the sequencer tracks (see section 21.2). Real-time operation is also dependant on the number of channels. This function is a way to run quick trials and decide whether a treatment is worthwhile.

To enter real-time mode, click on the button showing a loudspeaker, left of the Play button (keyboard shortcut: **Enter**).

To exit real-time mode, simply click once more on the button or press **Enter** again.

The same button is to be found in the dialog box that opens with Real-time processing Settings in the Processing menu (see section 28). Here you set the desired parameters. This allows you to listen to the result using different parameters.

In addition, in the Gain field, you can enter a an amplification factor value so as to avoid output saturation, since Normalization is not available in real-time mode.

After this verification, you start the processing in the usual way (Process Treatments, in the processing menu, please see section 25). The file will be written to hard disk.

29 Saving files

Any file type may be saved during or after the session. This is especially useful if you want to preserve files using different parameters. Saving is carried out via the File menu.

Only appropriate items appear in black, the others are grayed out. Certain items may be grayed out or may be different, depending on circumstances.

the File menu offers the following:

- **Save Treatments (Save Console, Save Bpf)** : This opens a standard dialog box allowing you to set the name and location of current files (or to save the console opened by New Console, or to save the Bpf).
- **Save Treatments As. . . (Save Console As. . . , Save Bpf As. . .)** : This opens a standard dialog box allowing you to set the name and location of current files (or to save the console opened by New Console, or to save the Bpf).
- **Save Sound As. . .** : opens a standard dialog box allowing you to choose the name and location of the current sound file.
- **Save Analysis As. . .** : offers a sub menu which in turn allows you to choose the analysis type to be saved, and then opens a standard dialog box allowing you to choose the name and the location where a SDIF file, containing an analysis, is to be saved.
- **Save Console. . .** : opens a standard dialog box allowing you to choose the name and location of the console attached to the AudioSculpt Window.

By default, saving is carried out in the preset folders contained in the application folder. You can change the default locations by defining new file paths in the Environment panel in Preferences.

For the My_Sound.aiff sound, the default name and folder are as follows:

- Treatments: My_Sound in the Treatments folder;
- Sonogram (Fft, LPC, Discrete Cepstrum); My_Sound.aiff.fft.sdif in the FFT folder;
- Markers: My_Sound.aiff.mrk.sdif in the Markers folder;
- Fundamental Frequency: My_Sound.aiff.f0.sdif in the Fundamental folder;
- Partials tracking: My_Sound.aiff.trc.sdif in the Fundamental folder;
- Chord sequence: My_Sound.aiff.mrk.sdif in the Fundamental folder;
- Peak detection: My_Sound.peak.txt in the Fundamental folder;
- Masking Effects: My_Sound.mask.txt in the SpectralEstimates folder,
- Processed sound: My_Sound.out.aiff then My_Sound.out 1.aiff. . . in the Sounds folder.

If you did not save, when closing the sound file a dialog box will invite you to save the desired components (the useful items will be checked, and you simply need to uncheck those you do not wish to save).

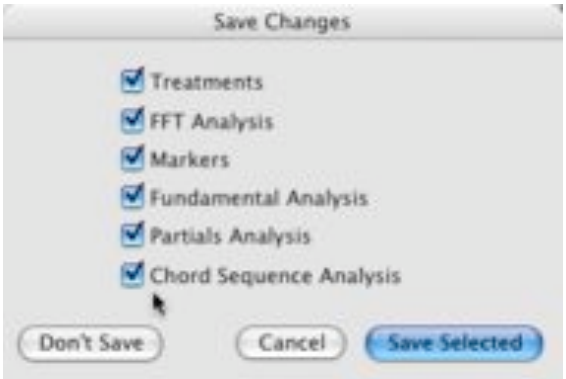

Note: If there is only one element to save, only the standard dialog box will open.

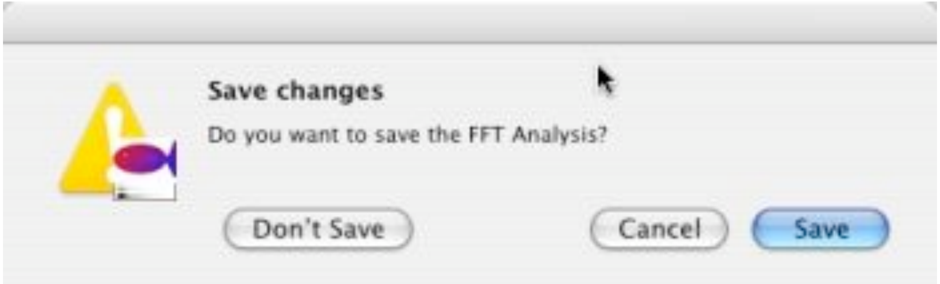

Treatments can also be saved using Drag and Drop onto the Desktop (see section 19.5.13).

30 File formats

30.1 The various formats

Processed sounds are saved under their initial format (AIFF, SDII or WAV), sample rate and quantisation resolution.

Treatments are saved under the SDIF format. Treatments saved via Drag and Drop are extracts in SDIF text file format. They can be read directly and are not modifiable.

Temporary files (treatment parameters) are text files, sometimes very long, and can be opened and edited in a text editor.

Consoles are in text format. This makes it easy to copy command lines for use in a SuperVp console.

Bpf's are saved in a format that is directly usable in Audiosculpt (they can be edited and changed in the Bpf editor). They are text files, sometimes very long, and can be opened and edited in a text editor.

> **Note:** This does not apply to analysis generated Bpf's.

Analysis are saved in SDIF or text format, according to type:

- **Sonogram Analysis, Fundamental Analysis, Partial Tracking Analysis and Chord Sequence Analysis** files are in SDIF format.
- **Peak Detection and Masking Effects** files are in text format.

30.2 SDIF files

The **SDIF** format (Sound Description Interchange Format) is an analysis and synthesis parameter file format, used for describing sounds.

For more information, consult the following site:

http://www.ircam.fr/sdif

These files can be imported into and used in other software (e.g. Open Music).

In AudioSculpt, by default, these files all have a double extension. For example: My_Sound.fft.sdif, My_Sound.f0.sdif or My_Sound.mrk.sdif. . .). The sdif extension must be preserved. It is advisable to preserve the other part of the extension in order to distinguish it from other files. Naturally, you can name it My_Soundfft.sdif or My_Sound_Sonogram.sdif or My_Sound_PrettyPicture.sdif.

An SDIF file can be converted to text and vice versa using Droplets. This makes it possible to edit them.

The Droplets folder, containing those Droplets (and the corresponding kernels) is in the AudioSculpt folder. Simply drag and drop the SDIF file onto the "sdiftotext" application and processing will start (there is a progress bar).

The My_Sound.aiff.fft.sdif file becomes My_Sound.aiff.fft.sdif.txt. After any editing, you can carry out the conversion in the other direction simply by dragging the text file onto the "text-

tosdif" application. The file becomes: My_Sound.aiff.fft.sdif.txt.sdif. This is the point at which it is advisable to simplify the name.

31 Menus

31.1 Menu File

Some items may be grayed out or may be different, depending on circumstances.

- **New Console** : opens a new SuperVp console, independently of any opened sound files. This allows you to start a processing series via command line.
- **New Bpf. . .** : opens a new Bpf editor window.
- **New Track** : adds a new track in the treatment sequencer.

- **Open. . .** : opens a standard dialog box allowing you to choose a sound file and open it.
- **Open Recent** : offers a sub menu allowing you to choose among recently opened found files.
- **Open Analysis. . .** : opens a standard dialog box allowing you to choose an analysis file as well as its associated sound, and open them.
- **Open Treatment. . .** : opens a standard dialog box for choosing a treatment file.
- **Open Analysis. . .** : opens a sub menu allowing you to choose the desired analysis type, and then a standard dialog box for choosing an analysis file (see figure 31.1-B).
- **Open Bpf. . .** : opens a standard dialog box for choosing a previously saved Bpf file.

- **Close** : closes the current window.
- **Save Treatments (Save Console, Save Bpf)** : This opens a standard dialog box allowing you to set the name and location of current files (or to save the console opened by New Console, or to save the Bpf).
- **Save Treatments As. . . (Save Console As. . . , Save Bpf As. . .)** : This opens a standard dialog box allowing you to set the name and location of current files (or to save the console opened by New Console, or to save the Bpf).
- **Save Sound As. . .** : opens a standard dialog box allowing you to choose the name and location of the current sound file.
- **Save Analysis As. . .** : offers a sub menu which in turn allows you to choose the analysis type to be saved, and then opens a standard dialog box allowing you to choose the name and the location where a SDIF file, containing an analysis, is to be saved (see figure 31.1-C).
- **Save Console. . .** : opens a standard dialog box allowing you to choose the name and location of the console attached to the AudioSculpt Window.

- **Page Setup. . .** : Not required for the Audiosculpt window; print format. . . for the console (opened via new Console) or for Bpf's.
- **Print. . .** : Not active for the Audiosculpt window; print format. . . for the console (opened via new Console) or for Bpf's.
- **Export Image. . .** : Allows you to export as desired, the waveform image of the entire sound, or of the sonogram, or of both, in PNG, TIFF, JPG, or PICT format (only active for the AudioSculpt window).

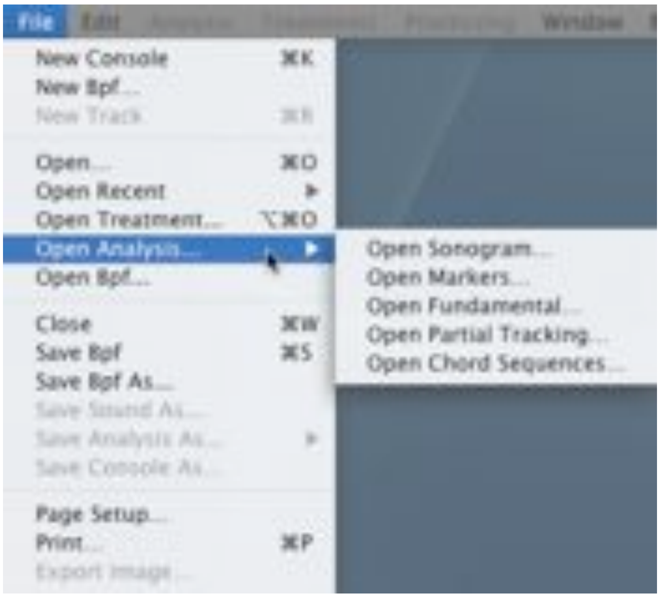

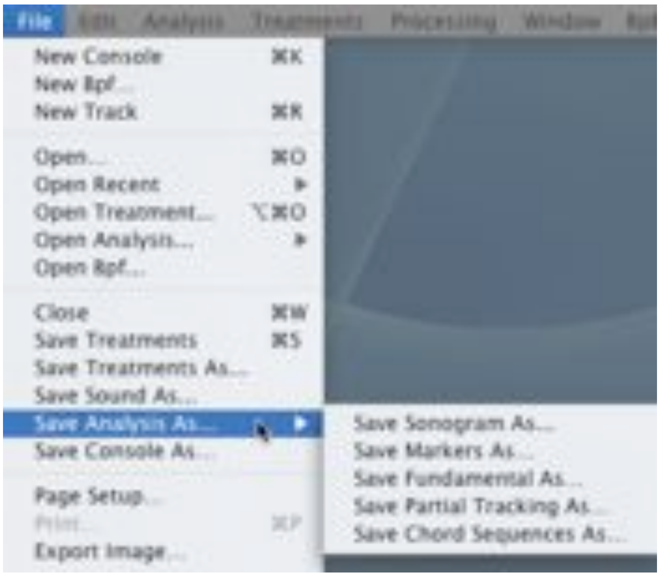

31.2 Edit Menu

Some items may be grayed out or may be different, depending on circumstances.

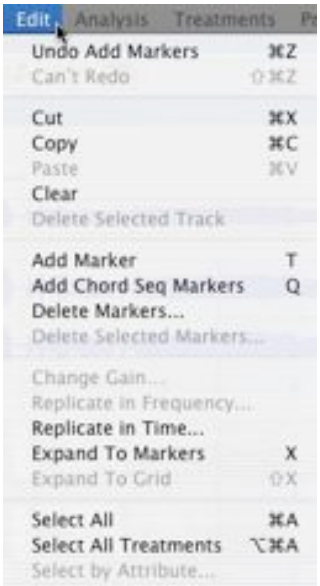

- **Undo** : cancels the last of certain actions (with a description of the action).
- **Redo** : redoes the last of certain actions (with a description of the action).

- **Cut** : deletes and copies into the clipboard (can be a treatment, a value, or a text).

- **Copy** : copies (can be a treatment, a value, or a text).
- **Paste** : pastes (can be a treatment, a value, or a text).
- **Clear** : deletes (can be a treatment, a value, or a text).
- **Delete Selected Track**: self explanatory.

- **Add Markers** : adds a manual marker.
- **Add Chord Seq Marker** : adds a marker pair for Chord Sequence Analysis.
- **Delete Markers...** : opens the Select the Markers To Delete dialog box, allowing you to choose the type of marker to be deleted.
- **Delete Selected Markers...** : self explanatory.

- **Change Gain...** : opens the Edit Gain dialog box that allows you to change the selected treatment gain (where applicable).
- **Replicate in Frequency...** : opens the Replicate in Frequency panel that allows you to adjust frequency duplication of selected surfaces (where applicable).
- **Replicate in Time...** : opens the Replicate in Time panel that allows you to adjust (positive or negative) duplication of selected treatment along the time axis.
- **Expand To Markers** : extends the duration of the selected treatments up to the nearest markers (left and right).
- **Expand To Grid** : extends the duration of the selected treatments up to the nearest grid references (left and right).

- **Select All** : selects the entire sound.
- **Select All Treaments** : self explanatory
- **Select by Attribute...** : not active.

31.3 Analysis Menu

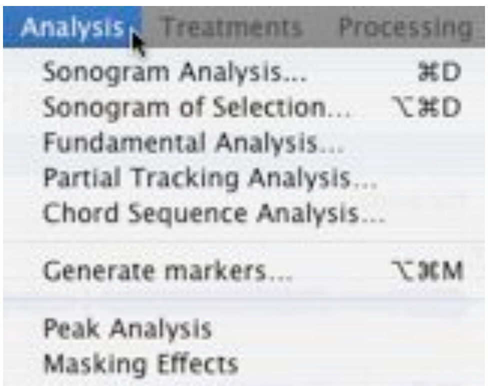

- **Sonogram Analysis...** : opens the Analysis Parameters panel for setting parameters and

starting sonogram display processing of the entire sound.

- **Sonogram of Selection...** : opens the Analysis Parameters panel for setting parameters and starting sonogram display processing of the selected part of the sound.
- **Fundamental Analysis...** : opens the Fundamental Analysis Parameters panel for setting parameters and starting analysis processing of the fundamental frequency (f0).
- **Partial Tracking Analysis...** : opens the Partial Tracking Parameters panel for setting parameters and starting analysis processing of partials.
- **Chord Sequence Analysis...** : opens the Chord Sequence Parameters panel for setting parameters and starting Chord Sequence processing.
- **Generate Markers...** : opens the Marker Parameters panel for setting parameters and starting non manual marker processing.
- **Peak Analysis** : opens the Peak Detection panel for setting parameters and starting peak detection processing.
- **Masking Effects** : opens the Masking Effects panel for setting parameters and starting masking effects processing.

31.4 Treatments Menu

By default, treatments will be applied to the entire sound. If only part of the sound is selected, treatments will be applied to that part only.

- **Add Constant Transposition...** : adds the corresponding treatment and opens the Constant Transposition panel that allows you to set Transposition parameters.

- **Add Constant TimeStretch. . .** : adds the corresponding treatment and opens the Constant TimeStretch panel that allows you to set parameters for stretching/compression.
- **Add Freeze. . .** : adds the corresponding treatment and opens the Freeze panel that allows you to set the freeze duration.
- **Add Reverse/Repeat. . .** : adds the corresponding treatment and opens the Reverse/Repeat panel that allows you to set treatment parameters.
- **Add Constant Formant Filter. . .** : adds the corresponding treatment and opens the Formant Filter panel that allows you to set formant filtering parameters.
- **Add Band Filter. . .** : adds the corresponding treatment and opens the Band Filter panel that allows you to set Band filtering parameters.
- **Add Clipping Filter. . .** : sets a Clipping filter that can be adjusted by the corresponding track element.
- **Add Image Filter. . .** : opens a standard dialog box allowing you to choose an image to be used as a filter.

For the following items, the Bpf editor opens a window that allows you to edit the corresponding Bpf.

- **Add Dynamic Transposition. . .** : adds the corresponding treatment and opens a window that allows you to edit a "Transposition" Bpf.
- **Add Dynamic TimeStretch. . .** : adds the corresponding treatment and opens a window that allows you to edit a Stretching/compression Bpf.
- **Add Breakpoint for Gain. . .** : adds the corresponding treatment and opens a window that allows you to edit a Gain Bpf.
- **Add Breakpoint Filter. . .** : adds the corresponding treatment and opens a window that allows you to edit a Filter Breakpoint.
- **Add Dynamic Formant. . .** : not active.

- **Invert** : allows you to reverse selected treatments (applicable to certain treatments only, in other cases, the item is grayed out).
- **Create Tracks When Needed** : When this item is checked, a new track is only automatically created if necessary (i.e. so that one treatment will not hide another).

31.5 Processing Menu

Some items may be grayed out, depending on circumstances.

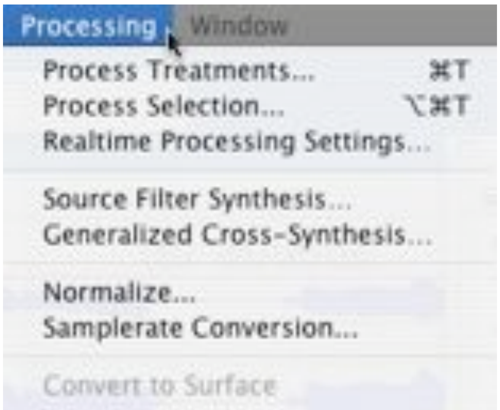

- **Process Treatments...** : opens the Processing Parameters panel for processing of treatments of the entire sound (remains grayed out until treatment is set).
- **Process Selection. . .** : opens the Processing Parameters panel for processing of treatments of the selected part of the sound (remains grayed out until treatment is set and part of the sound is selected).
- **Realtime Porcessing Settings. . .** : opens the Real-time Processing Parameters panel for real time mode.

- **Source Filter Synthesis. . .** : opens the Source Filter Synthesis panel that allows you to set source/filter type crossed synthesis.
- **Generalized Cross-Synthesis. . .** : opens the Generalized Cross-Synthesis panel that allows you to set general cross-synthesis parameters.

- **Normalize. . .** : opens the Normalize panel that allows you to set parameters for processing of this treatment over the entire sound.
- **Samplerate Conversion. . .** : opens the Samplerate Conversion panel that allows you to set parameters for processing of this treatment over the entire sound.

- **Convert to Surface** : not active.

31.6 Window Menu

- **Show Tools** : displays the floating Toolbox (if the option is checked).
- **Show Inspector** : displays the floating Inspector (if the option is checked).
- **Show Sonogram Display** : displays the Sonogram toolbox, which allows you to:
  - show/hide zone 3 elements (sonogram).
  - adjust the sonogram grayscale levels, i.e. black and white sensitivity levels.
  - to adjust the marker display threshold.
- **Show Grid Settings** : displays the Grid Settings toolbox.

  Note: All toolboxes (palettes) that are checked will open automatically in the following session. The **tab** key allows you to show/hide these toolboxes (Tools, Inspector, Sonogram Display and Grid Settings).
- **Maximize Sonogram** : allows full screen display of the sonogram and the "window into the sound" (zones 2 and 3) or return to the initial display. The keyboard shortcut is: **Apple+U** (**Command+U**).
- **Follow Playback** : causes the window to scroll the sound along with the position indicator pointer [line].

At the bottom is a list of open windows. It allows you to bring any one of them to the foreground. Keyboard shortcuts make rapid operation easy.

31.7 Bpf Menu

This menu is only displayed when the Bpf editor window is open.

- **Show Bpf Tools** : opens the tabbed Bpf window (it can also be accessed directly via the Bpf window).
- **Show Points** : displays a Bpf's points (dots).
- **Show Double Line** : when checked, the Bpf is displayed with a thicker line.
- **Resampling. . .** : resamples the Bpf.

The other commands (items) are grayed out and do not apply to AudioSculpt.

32 Command lines

For advanced users, SuperVp processing can be started directly from AudioSculpt using command lines. Simply type in (or paste, or drag and drop) a line into the currently active console and press **Enter**. Processing will begin.

Create new consoles using the New Console command in the File menu. You can use parameter files (especially those generated by AudioSculpt (see section 25.3).

The command line must be written as a single line, without carriage return. You can of course type more than one line, they will wrap down automatically.

> **Note:** This type of processing calls SuperVp directly. Therefore, the sounds do not open automatically (they are stored in the locations you defined). When carrying out a series of treatments, do not forget to give each file a different name, because they are not automatically renamed.

The console can give you very valuable information on remarks or error messages returned by SuperVp. AudioSculpt will mostly display the essential information.

You can find the following help files: SuperVP "SuperVP - svphelp" and "Filter module description" in the Appendix (section 33) as well as attached text files. The console supported by AudioSculpt cannot display them fully.

You can also launch SuperVp via the "Terminal". If you do, it is advisable to use the disk image version on the CD-Rom (SuperVP.dmg).

33 Annexes

33.1 SuperVP -ha

```
===================================================
SuperVP (IRCAM) 1990-2005       version : 2.59j fix-1
(compiled by niels  for AudioSculpt
on Wed Jan 26 17:45:54 GMT 2005) Prec=double
===================================================

helpoptions::
==============

   -h   :prints out this help description.
   -ha  :prints out all sections of the help besides the extended
            filter description.
   -hi  :prints out the help message for input options.
   -hp  :prints out the help message for processing options.
   -hf  :prints out the description of the different filtering modules.
   -ho  :prints out the help message for output options.

Generally, SuperVP has two input tracks, if applicable on both tracks
the options for track 1 are in upper case, for track 2 in lower case.

Option flags specified in <> are mandatory. If option flags are
specified in [] they may be omitted resulting in a default value.

Parameter files: The parameter files that are not referenced by means of
a complete path name (either starting with "/", "./", or "../")
will be prefixed by the directory stored in the environment variable
SVPPARAMS! If SVPPARAMS is not set the current directory is used.

----------------------------------------------------------

Input options ::
================

-Ss<filename> : specifies input file name (Def: stdin)
                which is relative to $SFDIR  if no full path is given
                filename can be either a soundfile or a data file
                (see output options -Og3 or -Og5 format)
                generated by an earlier call to SuperVP.
                To specify that a data file is used the corresponding
                track should be switched to data input mode with the
                -isdata flag.


                Supported sound files comprise AIFF/AIFC/WAV/NEXT/SDII
                with all variations of sample size and with u-law/a-law
```

```
                compression.

 ------ Example: -Sflute

                Specifies input data to be read from $SFDIR/flute!
----------------------------------------------------------

-isdata <track>: File specified for given track (S for first track,
                 s for second track) contains spectral data
                 in g3 or g5 format (see input options -S and
                 output options -Og3 or -Og5 format)

 ------ Example: -Sflute.g3 -isdata S

                Specifies input data to be read from $SFDIR/flute.g3
                is in gabarit format
----------------------------------------------------------

-Bb<start> :     specifies start position (Def: 0.0) in the input file.
                The starting point may be negative which will
                add silence in the beginning of the sound.

                If start contains a decimal point it is interpreted as
                time in seconds if not it specifies time in sample number!
                The first sample is at position 0;

 ------ Examples: -B1.4    starts processing 1.4s after the first sample
                  -B44100  starts processing at sample position 44100
----------------------------------------------------------

-Cc[k]<channel> :  specifies the channel to process (Def. all channels or 1).
                Channel numbers start from 1. The default value depends
                on the selected output format. For sounds the default
                is to process all channels, for analysis data only one
                channel can be processed and the default channel is channel 1.

                The modifier k is only supported for the first track '-C'
                and is used for multi channel files to
                keep the channels not selected for processing unchanged
                in the output file.


 ------ Examples: -C2   selects the second channel from a multi channel sound file.
                        and creates a mono output file.
                  -Ck1  selects the first channel from a multi channel sound file
                        and creates a sound file with the same number of channels
                        then the input and only the first channel processed.
----------------------------------------------------------

-Ee<end> :       specifies last selected sample in the input file
                (Def: end of file), if end points to samples past the end of
                file zeros are added to the file content.
```

```
                If end contains a decimal point it is interpreted as
                time in seconds if not it specifies time in sample number!

 ------ Examples: -E5.0
                  -E88200
----------------------------------------------------------

-inplace : stores results into input file. Especially useful when
        processing selected segments (-B/-E) in which case
        only part of the sound is processed and the result is put back into
        the input file.

        Normalization is not available during inplace processing.
        When inplace processing is used with time dilation (-D)
        the selection of the end of the segment should be carefully
        selected such that the phase shift due to time dilation will not
        lead to artifacts. Therefore a low amplitude region should
        be used.

    ------ Example: supervp -Sfile.aiff -inplace -Fbande-noex band.par -B1. -E2.
        Will apply the band filter with parameters band.par to the sound
        file.aiff. Only the segment from time 1 sec up to 2 sec
        is processed and the result is put back into  the original sound.

---------------------------------------------------------------------

Processing options ::
=====================

-Aa[analysis type (Def: fft)] [] : specifies the analysis procedure. type is:

    fft [order (Def: 30)] : fast Fourier transform
    ------ Example: -Afft
           if transposition with preservation of spectral envelope is performed
           the lpc order for the estimation of the spectral envelope
           is given by means  of the additional parameter order.
    ------ Example: -Afft 25
           uses lpc order 25 for the preservation of the spectral
           envelope. Note, that the spectral envelope will only be preserved
           if transposition has been specified by means of -transke.
    lpc [order (Def: 30)] : linear prediction analysis. order is the
           number of poles. The lpc filter is multiplied by the
           residual energy to have a proper spectral representation.
           See -ns/ -truelpc switches for further information on lpc
           normalization.
    ------ Example: -Alpc 50
    lpc_inv [order (Def: 30)] linear prediction analysis with
           envelope inversion. order is the number of poles.
           LPC analysis with with inverted spectrum for inverse
           filtering (for possible application see -G option or -Fgabarit).
           The inverse lpc filter is multiplied by the inverse
```

```
       residual energy such that it may be used to create
       normalized excitation signal with the -Gmul mixer.
       If you apply inverse lpc filtering using the same signal
       for both tracks you create an energy normalized
       excitation signal that may be used as excitation for an
       lpc spectral envelope obtained with a second
       -alpc analysis recreating the energy contour of the related
       second signal.
       See -ns/ -truelpc switches for further information
       on lpc normalization.
------ Example : -Alpc_inv 50
freeze [nbframes (Def: 30)]
       Collect frequency statistics for nbframes and then freeze
       the amplitudes at frame nbframe and modulate the phase
       of the synthesized frequencies according the frequency
       tables gathered so far.
------ Example : -Afreeze 42
newfreeze [parfile|parstring]
       freeze the sound at the given time instances for the
       specified duration. Multiple freeze points can be established
       with the newfreeze modul.

       Parameter string for each freeze point defines

        start, duration, stat_length

       where
        start       specifies the start time of the freeze
        duration    specifies the duration of the freeze
        stat_length specifies the duration over which frequency
                    statistics for each frequency bin
                    are gathered before the freeze

       All values are specified in seconds.
       The parameters may be specified on the command line as a sequence
       of comma separated floating pint values or in a parameter file
       as a sequence of white space separated floating point values.

       After each freeze the original sound is
       continued. If the last character of the parameter string
       is the character e then after last freeze period the sound
       simply decays to zero without continuing
       with the original sound.
       The frequency statistics tables can collect a maximum
       of 100 frames. Time ranges that extend over more frames are
       simply reduced to fit this limit.
       The  frequency statistics are used to modulate the
       phases of the synthesized frames. More variation in frequency
       will result in increased phase modulation which in turn
       increases amplitude modulation which may reduce the
       synthetic impression of freezed sounds.
```

```
       The example below freezes the input sound at time
       0.1s for 0.4s without using the frequency statistics
       and at time 0.5s for 0.6s modulating phases
       according to the frequency variations seen over the last 0.1s
       The sound ends after the second freeze by simply decaying to zero
       within half the size of the analysis window. For the same
       result the freezefile would contain the same string
       however, replacing the commas by white spaces.
       The newfreeze module supports phase synchronization see -P.
------ Examples : -Anewfreeze 0.1,0.4,0.,0.5,0.6,0.1,e
                  -Anewfreeze freezefile
pic [theshold] [nnumber]: peaks detection
       threshold is in dB attenuation below
           the peak amplitude, (Def: no threshold)
           a peak is only taken if its maximum is at least threshold dB
           above the neighboring minim.
       number is the max. number of peaks in the output, the peaks selected are
           the  peaks with the largest amplitude. For output they
           are sorted according to frequency
       (attention: no blank between n and value)
------ Example: -Apic 30 n20
        detect 20 strongest peaks that have amplitude larger than 30db
        above the two neighboring minima.
mask [list] [threshold] [nnumber]: Terhardt algorithm for
       spectral smoothing list is a list of words in:
       amp, freq, weight, all, (def: weight)
       threshold (see -Apic)
       number is the max. number of peaks, for mask peaks are selected according
       to their weight.
       (no blank between n and value)
------ Example: -Amask amp, weight 10 n20
ced [order] : discrete cepstrum analysis. Provides a spectral envelope.
       default order is (Def: 30).
------ Example: -Aced 55

f0 [fmmin,][fMmax,][FFmax,][snthreshold,][smoothorder] : pitch detection
       fm is followed by minimum value for pitch in Hertz (Def: 50)
       fM is followed by maximum value for pitch in Hertz (Def: 2500)
       F  is followed by maximum frequency in spectrum (Def: 7500)
       sn is followed by noise threshold in dB (Def: 50)
       smooth is followed by median smoothing order (Def: 1 (no smoothing))
             smoothing order is enforced (rounded up) to be uneven!

       There is no blank between the parameter name and its value.
       The output file will be in SDIF format as long as -Oa is not given
------ Examples: -Af0 "fm100, fM1000, F2500, sn60, smooth3"
                     -Af0 fm10,sn25
formant_lpc [nnumber] [lpc order] : formant detection
       from a spectral envelope computed via lpc.
       In this analysis a formant is understood to be a peak of
       the spectral envelope calculated using the lpc analysis module.
       number is the max. number of formants to detect (Def: 5)
```

```
          (ordered according to increasing frequency) (no blank between n and value)
          order is the lpc order (Def: 15).
   ------ Example: -Aformat_lpc n3 10
   formant_ced [nnumber] [ced order] : formant detection
          from a spectral envelope computed with discrete cepstrum.
          In this analysis a formant is understood to be a peak of
          the spectral envelope calculated using the ced analysis module.
          number is the max number of formants to detect (Def: 5)
          (ordered according to increasing frequency) (no blank between n and value)
          order is the ced order (Def: 45).
   ------ Example: -Aformant_ced n16 55

   ced_inv [order] discrete cepstrum analysis with envelope inversion.
          Use it like CED analysis but only for inverse
          filtering (see -G option).
          Order is the ced order  (Def: 30).
  ------ Example : -Aced_inv 50

Please read -O section for the formatting of the output data.
----------------------------------------------------------

-avseg <file or pair of time values> :
        computes averaged spectra for the time segments specified.
        If a file name is given this file should contain pairs of start - and
        end times of the segments. The segments can also given as a comma
        separated list of times directly on the command line.
        If only a single number is given single segment is formed that
        starts at the time given and continues until the end of the sound.
        The segment positions need to be in increasing
        time order and cannot overlap.

        You need to select the output options
        -OM1 or -OM2 to obtain the spectral
        averaged output. See -OM for further information
        on the output format.

        ------ Example: -avseg 1.0,1.1 -OM1

        adds an segment lying between 1.0s and 1.1s to the list of segments.
        and stores the averaged spectrum in the output file.
----------------------------------------------------------

-D[coefficient (Def.: 1)] or -D<filename> : constant or time-varying
                                                         time-stretching

 Apply time dilation to input sound.

---   Examples: -D2.0  (dilation of 2, duration is doubled)
                -D0.5  (compression, duration will be divided by 2)

   parameter files are used for time varying operations. The files contain
   multiple lines indicating:
```

```
      time dilate_coefficient

    The entries are linearly interpolated to derive the time dependent
    dilation. The dilation coefficient before the first and after the
    last file entry is extrapolated keeping its value constant.
    The -D flag may be given more than once in which case
    the result is a multiplication of the requested individual values.

    Time varying dilation and transposition may be applied in a
    single SuperVP call in which case the requested dilation is and the time
    compensation for the transposition are added.

    The optimal analysis step size is calculated automatically if no -I option
    is specified. This is recommended especially for time-varying
    dilation/transposition.

    You may use the -P switch for phase synchronized processing.

----------------------------------------------------------


-Ff<filtertype> <filename> specifies filter
    filtertype is one of
       bande
       gabarit
       breakpt
       surface
       fof
       fifof
       clip
       denois
       fshift

    filename is a file containing filter parameters.

    Multiple filters can be applied during one call to SuperVP.
    use option "-hf" for a description of available filters
    and the related parameter file format.

 ------ Example: -Fbande file
----------------------------------------------------------
-Ff<superposition mode>

    If multiple filters of the types bande, gabarit, breakpt,
    surface, fof, fifof are applied the standard behavior superimposes
    the filters by means of multiplying the transfer functions.
    This behavior is explicitely selected by specifying -FCombineMul.
    The alternative is -FCombineMax which will only apply the
    filter with maximum amplitude.

    The Mode of Superposition may be selected independently
    for the two tracks
```

```
 ------ Example: -FCombineMax -fCombineMul
----------------------------------------------------------

-G<cross-synthesis type> [filename] or
  <cross-synthesis type> [-X<val>(Def:1.)] [-x<val>(Def:1.)]
                     [-Y<val>(Def:1.)] [-y<val>(Def:1.)] [-q<val>(Def:1.)]

   Combine the two different tracks into a single track. Suppose S1 S2
   are the short time FT spectra of the two input channels and SO
   is the resulting spectrum of the  cross-synthesis and X/x/Y/y/qval is
   the value specified with flag -X/x/Y/y/q then the mixer operation
   is for

   cross type add: transform input spectra into representation
        using real/imaginary part and calculate

          SO = Xval * S1 + xval * S2,

   cross type cross: S1 and S2 are calculated in amplitude/freq
          representation and the output is

          Amp(SO) = Xval*Amp(S1) + xval*Amp(S2) +qval*Amp(S1)*Amp(S2)
          Fre(SO) = Yval*Fre(S1) + yval*Fre(S2),

   cross type mul: (source-filter), S1 and S2 are transformed
        into Amp/Phase representation and the output is

             Amp(SO) = Amp(S1)*Amp(S2)
             Pha(SO) = Pha(S1)+Pha(S2),

   cross type amul: (source-filter), S1 and S2 are transformed
        into Amp/Phase representation and the output is

             Amp(SO) = Amp(S1)*Amp(S2)
             Pha(SO) = Pha(S1),

   cross type pmul: (source-filter), S1 and S2 are transformed into
        Amp/Phase representation and the output is

             Amp(SO) = Amp(S1)
             Pha(SO) = Pha(S1)+Pha(S2).

   filename is the parameter file for add and cross modes.

        for cross mode the file is made of lines with: time X x Y y q
        for add mode the file is made of lines with: time X x

 Notes:
   Order and position of the -XxYyq flags is free.

   The processing is stopped in cross and add mode if both files are finished,
```

```
   in mul/amul/pmul mode if either source is finished.

   Note that the increment step of the second track is adapted
   to match the duration of both sources if you do not specify the
   step size for the second channel with -i explicitely

------ Examples: -Gadd file
                    -Gadd -X1.0 -x2.0
                    -Gcross -X0.5 -x0.5 -Y1.0 -y1.0 -q0.0

  with mul mode inverse filtering to obtain a normalized excitation
  can be performed with the same  sound on both channels.
  ------ Example :
     supervp -Ssound_file -ssound_file -Gmul -Z -Afft -alpc_inv 15 output_file
----------------------------------------------------------

-ggain <gain_factor or filename> or -gtremolo <mode> <filename>
            -gfilter <filter_parameters or filename>: amplitude modulation
   gain <filename>: multiplies output samples with envelope specified in file
                    file contains lines with time envelope
   tremolo <mode> <filename>
           mode is the type of amplitude modulation
            sinus
            carre (square wave)
            triangle
            scie (sawtooth wave)
            file contains lines with time depth
   filter <parameters or file>
        filter signal using a linear filter that implements
        o(n) = sum_i b_i x(n-i*K) - sum_j a_j o(n-j*L).
        the filter parameters b_i,a_j are either specified directly
        on the command line or in the parameter file.
        The format is : K,b_0,b_1,...;L,a_0,a_1,..
        If a_0 is not equal 1 all a_i are normalized by a_0.
        The semicolon can be replaced by means of a slash /

        ------ Examples: -ggain file
                         -gtremolo sinus file
                         -gfilter 1,1,0.2;3,1,0.5
----------------------------------------------------------

-H<sample rate> : specifies new sampling rate
   the header of the sound file will be changed
   If this option is used, don't use -A and -Z.
        ------ Example: -H32000
----------------------------------------------------------

-Ii<step> or -Ii<method> <filename> or -I<num/mem> <step>
   or Ii<pos> filename step:
    specifies analysis step
    the step for analysis is in samples
    (Def I : max(1,min(window_size/8,window_size/time_dilation/8))
```

```
(Def i : step size track 1 /*length track 2/length track 1)

<method> is the method for reading parameter file name
    num  same as giving a numerical value (no parameter file)
    mem  input is callback (only for use in svp library applications)
    sync ordered pairs in parameter file (time, fundamental
         frequency) (not working)
    dep  ordered pairs in parameter file (time, step value) (not working)
    pos  sample position in the sound file is given for each
         window, the additional step will be used as virtual step (vstep)
         for the analysis and resynthesis. This parameter is only needed
         when resynthesis (-Z) is requested. The file consists of
         lines containing:

         single position       : indicating the  sample
         (ex: 10000.)            position to read the next frame
                                 The frame is positioned in
                                 the output stream with
                                 location given by vstep.
         "c" followed by pos   : indicating a repositioning
         (ex: c 10000.)          of the frame position
                                 without really creating any output.
         "u" followed by pos   : indicating a target position
         (ex: u 10000.)          The frame increment is calculated
                                 automatically and a sequence
                                 of frames is created that stops with the last
                                 frame having its center at
                                 position "pos".
         "f"                   : automatically move forward until
                                 a new position can be read from
                                 the parameter file, only useful
                                 if reading from a named pipe.
                                 Position increment is the same
                                 as with "u pos"
         "b"                   : automatically move backward until
                                 a new position can be read from
                                 the parameter file, only useful
                                 if reading from a named pipe.
                                 Position increment is the same
                                 as with "u pos"

    ------ Examples: -I256
                     -Ipos posfile

        example posfile:
        u 30000            (: read frames until sample position 30000)
        c 100000           (: position frame at sample position 100000)
        u 30000            (: read frames backward until sample pos 30000)
        c 100000           (: position frame at sample position 100000)
        u 200000           (: read frames until sample position 200000)

        The sequence of position commands will process the file
```

```
            and apply all treatments using the sound segment between
            sample 0 and sample 200000 with the special effect to
            reverse the sound samples located between sample 30000 and 100000.

----------------------------------------------------------

-J<type> : specifies synthesis window type
        (Def:same as the principal track)
        supported windows see -W
----------------------------------------------------------

-logfile filename : sets message output file to filename, the filename
stderr is treated as special indicator for stderr output
----------------------------------------------------------

-Mm<window size> or -M<method> <filename> :
                          specifies analysis window size
    the window size for analysis is in samples (Def: 1024)
    <method> is the method for reading the parameter file name
        num  same as giving a numerical value (no parameter file)
        sync ordered pairs in parameter file
                      (time, fundamental frequency) !!currently not supported!!
        dep  ordered pairs in parameter file
                      (time, window size) !!currently not supported!!
        ------ Examples: -M2000
                         -Mdep file
----------------------------------------------------------

-N<FFT size> : specifies FFT size
    the FFT size in samples (Def: 1024)
    must be greater than window size, see -Mm
        ------ Examples: -N4096
                         -N8192
----------------------------------------------------------

-nn do not normalize output sound file (this is default now)!
   ATTENTION : The use of -nn is strongly discouraged. Since version 1.75
   no normalization is the default behavior and -nn flag is no
   longer supported and will be switched of in future versions!
   See -norm flag
---------------------------------------------------------------------

-norm [level (Def.: 0)] normalize output sound file to level dB below
   maximum range of the output data type.
   0dB normalization produces data in the range
   between +/- 1.0  for floating point sample format and
   in the range between +/- (2^(N-1)-1 for N Bit integer sample format.

   NORMALIZATION Behavior changed!! If -norm option is not given normalization
   is NOT performed.

   IMPORTANT : when normalizing, SuperVP uses a temporary file. It's length is
```

```
   the same length of the processed sound file when using 32bit floating point
   packing mode or twice the length when using 16bit short packing mode.
   Temporary file is created in SVPTMP (setenv SVPTMP mytemporarydirectory).
   When SVPTMP is not set, SFDIR is used.  If SFDIR is not set, current
   directory is used. If the chosen directory is not writable, then /var/tmp
   is used.
----------------------------------------------------------------------

-ns requests spectral normalization for analysis output, which
   results in amplitude values to be below 1 (Def: no normalization)
-----------------------------------------------------------

-oversamp <val> :
        specfies the minimum overlap of analysis and synthesis windows.
        The maximum increment between successive windows
        is specified in terms of parts of the window size (Def: 8).
        Smaller values indicate less overlap which results in less
        quality but also reduced computational costs.
        The default value (increment not larger then an 8-th part of the window)
        ensures very high quality for all situations. Values above  4
        will be almost always sufficent. For values below 4
        artifacts are unlikely but may occur. Minimum value is 2.

        ------ Examples: -oversamp 4


-----------------------------------------------------------

-P[tp] : switches on phase synchronization (Def: no phase sync)
      and possibly transient preservation
      mode for time dilation (see -D) or newfreeze modul.
      It generally improves amplitude reproduction and decreases
      the phasiness when processing non stationary sounds. This switch
      will be ignored if no dilation/newfreeze is requested.

      Transient preservation is switched on if the
      tp is larger than 0 (Def: no transient preservation/tp=0).
      Transient preservation localizes transients peaks in
      time/frequency domain and re-initializes phases for the related bins
      after the transient has passed in order to keep
      the wave form of the transient.

      transient detection parameters can be changed by means of
      -td_thresh/-td_G/-td_nument/-td_band/...

-----------------------------------------------------------

-Rr<samplerate> : specifies sampling rate
    used when performing resynthesis on analysis data file
-----------------------------------------------------------

-resS <value> : specifies stopband attenuation of the interpolation filter
```

```
        used for resampling (Def: 70dB)!
        The window size and oversampling of the interpolation filter
        are automatically adjusted to achieve the requested
        attenuation of aliasing resulting from the interpolation
        process! To prevent excessive filter sizes attenuation of
        more than 140dB are not recommended.

----------------------------------------------------------

-t  :displays the current processing time in the input file
----------------------------------------------------------

-T  :displays the current processing time in the output file
----------------------------------------------------------

-td_int  :
         switch to require integration of mean time for transient
         detection over td_nument peaks. This will improve
         robustness against detection of noise (Def: 2)

        ---- Example: -td_int
----------------------------------------------------------

-td_thresh <value> :
        adjust threshold of peak mean time (center of gravity of the energy
        of the signal related to the peak)      to detect a transient peak
        (Def: 1.4). The mean time to end a transient is
        related to the mean time of a simple ramp covered by the
        analysis window. The amount of access mean time to detect a start
        of a transient is determined by td_thresh. By increasing
        td_thresh the transients need to be more pronounced
        with respect to the noise background to be detected as
        transient. The parameter is required to be above 1
        because other wise the transient would be ended before it began.
        Optimal thresholds depend on the noise in the signal and on the
        confidence threshold determined by -td_G, however, a range between
        1.2 2.5 is usually reasonable.

        ---- Example: -td_thresh 2

----------------------------------------------------------

-td_G <value> :
        confidence factor to use when comparing the number of attack
        transient peaks in the transient statistic
        bands current frame and comparing it to
        the number of attack transients in the related band
        in the previous frame.
        The larger the value the higher the number of attack transient
        peaks in the current frame needs to be to be
        detected as part of an attack transient.
        The value denotes the access of the  transient peak
```

```
        frequency estimated in previous frames in terms
        of the  standard deviation (Def: 2.5)


        ---- Example: -td_G 2

        Reduce transient attack confidence by requesting
        only an access of 2 times the standard deviation
        for a band to be detected as transient.


----------------------------------------------------------

-td_band <value,value> :
        band  in Hz used for transient detection.
        only spectral peaks within the band will be used to
        determine transient regions. (Def: 0,sample rate/2)

        See: -td_nument

----------------------------------------------------------

-td_nument <value> :

        band size  used for the statistical monitoring
        of noise related background transient activity.
        Band size is specified in terms a number of peaks (mainlobe width)
        that fit into the band (Def: 10).

        This parameter controls the impact of a single transient peak
        for the detector. The smaller the number of peaks in the
        band the more impact a single sinusoid has. The allows to
        detect transients more sensitively even if they are close
        to stationary signals. At the same time it will increase the
        probability of false detections in noisy regions.

        ---- Example: -td_nument 30

        requires to collect transient statistics in bands that may
        hold 30 stationary sinusoidal peaks.

        See: -td_G

-----------------------------------------------------------

-td_ampfac <value> :
        specifies factor to use to compensate the missing
        contributions from previous frames when restarting
        after a attack transient (Def: 1.8).

        ---- Example: -td_ampfac 1.
```

```
        when restarting phases after attack transient multiply
        amplitude by 1. instead of default value 1.8.
        This will remove the compensation of the missing amplitude
        info from the frames before the attack transient has been
        restarted.

----------------------------------------------------------

-td_mina <value> :
        specifies minimum amplitude that an attack transient
        needs to achieve to be detected (Def: 0). The amplitude
        reference is interpreted as a normalized value.

        ---- Example: -td_mina 0.01

----------------------------------------------------------

-td_nosync  :
        requires to not syncronize transients detected in different
        channels of a multi channel audio file. Per default the
        transients of the different channels are synchronized to
        prevent the artifacts resulting from the fact that the
        transients in the different channels may be detected at different
        places or may be detected only in one of the channels.

        ---- Example: -td_nosync

----------------------------------------------------------

-td_declick <value> :
        specifies maximum duration of click events in seconds.
        All transients that start sound events of duration shorter
        than the given duration will be removed from the signal.

        ---- Example: -td_declick 6.3492e-04

        Removes clicks that have been detected as transients
        with duration shorter than 0.63492ms.

        Note that click removal is related to transient
        detection and therefore it will only work if
        the -D (with any argument) and the -P1 flag have been
        given on the command line.

----------------------------------------------------------

-trans <cents> or <filename> : transposes by the given number of cents
        applying time correction to compensate the
        duration effects of transposition by means of resampling.
        Transposition requires the -A and -Z flag.
        For time varying transposition a parameter file has to be used
        which contains lines with:
```

```
                time transposition

            Which are linearly interpolated. Since supervp 2.40 the trans flag
            may be given more than once in which case the effective
            transpositions are added. transpositions are always
            extrapolated outside the range specified in the parameter file.

        ------ Example: -trans 1200 (will transpose one octave up)

    -transnc <cents> or <filename> : Applies transposition without
            time correction. The perceived pitch is changed by means of
            sample rate conversion without changing the samplerate
            of the resulting signal. Therefore the length of the sound
            will be changed. The parameters or parameter  files and
            are equal to the -trans flag. -trans and -transnc flags may
            be given more than once even with overlap in which case
            the requested transpositions are added but will only be
            partly compensated.

        ------ Example: -transnc 1200 (will transpose one octave up)

    ----------------------------------------------------------

    -truelpc : do not apply any normalization after lpc in the analysis section.
               Intended for using lpc analysis for lpc filtering.
    ----------------------------------------------------------

    -Uu : specifies the position of the first window with respect to
          the first sample (Def: window centered at first sample)

          -U/-u moves the window such that its first sample matches the
          first sample of the sound.
    ----------------------------------------------------------

    -v  :prints out details about the SVP patch structure and parameters
    ----------------------------------------------------------

    -Ww<type> : specifies analysis window type
        the analysis window type (Def: hanning)
        supported windows:
           rectangular (rect, rectangle)
           triangular  (triangle)
           hamming
           hanning
           blackman
           exactblackman

     ------ Example: -Wblackman
    ----------------------------------------------------------

    -Z : performs a resynthesis (inverse FFT and overlap/add)
```

```
-----------------------------------------------------------

Output options ::
=================

<output filename>
  The outputfile is specified without option switch as
  last parameter of the command line.
  If no output file is specified or the name of the file
  is "stdout" the result is directed to stdout.
  if the name of the output file is "@PLAY@" the output is not
  stored as a file but directly played as sound via the audio hardware
  The default latency is 0.5 sec and can be changed by means of adding a
  float number to the play file name. For example @PLAY@0.2 would indicate
  play the sound with latency 0.2 sec

-O<mode>:<list> or -O<mode> <file>
  <mode> is the output type (b)
  SOUND modes :
   sa  for AIFF/AIFC  16bit integer soundfile format.
       sA8,sA16,sA24,sA32 select the different sample size in bits
       sAf,sAd  select float or double samples in AIFC format.
   sA  equivalent to Sa
   sis for Ircam 16bit integer soundfile format.
   sif for Ircam 32bit float   soundfile format.
   sw  for wav 16bit integer sound file format.
       sw8,sw16,sw24,sw32  select the different sample size in bits
       swf,swd select float or double samples in WAV format.
   sW  equivalent to sw.
   sn  for NeXT  16bit integer soundfile format.

ss  for Mac Sounddesigner II format

srs for raw   16bit integer soundfile.
   srf for raw   32bit float   soundfile.

  IMPORTANT : default format and packing mode (integer or float) differ
   according to the output type (a file or a pipe):
        File output : those of the input processed sound (not normalized).
        Pipe output : raw 32bit float (not normalized).

ANALYSIS Output Modes :
   a for ASCII analysis data output
   b for binary analysis data output (default)
   g0             : Unified file format, data is log amplitude
                    stored as unsigned char
   g1,g2,g3,g4,g5: Unified file format, see below and gabarit filter
   S0,S1,S2,S3,S4,S5: Unified file format stored as SDIF file
   T              : sdif file containing description of attack transients
   F+,F-          : sdif file containing local maxima
                    of averaged spectral differences
                    (add selects regions with  increasing (+)
```

```
                    or decreasing (-) amplitude for
                    difference evaluation, Def: +)

   M1,M2:          : averaged spectral information for all segments specified
                    via -avseg flag. Output is stored in an SDIF file in an 1AVS/1AVS
                    Matrix following an IAVS frame denoting analysis infos.
                    The M1 mode  outputs the mean and standard deviation
                    of the absolute spectrum for each bin, the M2 mode
                    outputs maximum and standard dev. of the absolute
                    spectrum for each bin. Note, that the power spectrum
                    can be derived from M1 results by (mean*mean + stddev*stddev)!

Format ( binary file + header like gabarit file ) using
   FFT, LPC, CEPSTRE.
   File content is respectively t,amp ; t,phase ; t,amp,phase ;
   t,freq,amp ; t,pr,pi.
   g1,g2 formats work with gabarit filtering and g4,g5 ones
   work with the synthesis module with data on input channel.
   g3 format works with both.
   Use these format with -Oxx option (no more arguments on
   command line)
   ------------- Ex : -Og3
   Using other format arguments will print a warning but
   the right output format will be set.

   <list> is a list of strings separated by ','
         number, time, amplitude, frequency, phase, etc..
   <file> is a text file where output format is specified
    ------ Examples: -Oa:number,amplitude,phase
                         -Oa foo.format
----------------------------------------------------------

List of available strings as output data
        t,time: time (of frame)
        number: number (of entries per frame)
        a,amp,ampl,amplitude: linear amplitude
        adb,amp_db,ampl_db,amplitude_db: amplitude in dB
        pr,partie_reelle: real part of a complex number
        pi,partie_imaginaire: imaginary part of a complex number
        f,fhz,freq,frequence: frequency in Hz
        midicents,fmc: frequency in midicents
        i,ind,index: frequency as fft-index
        chan,chan_freq: center frequency of fft-channel
        ph,phase: phase
        spl: sound press excess
        weight: harmonic weight
        tsp,truesp,truepitch: harmonic true pitch
        sc,score: score
        coeff,coeff_pitch: pitch confidence
        largeur,largeur_formant: formant width
        cf,coeff_filtre: autoregressive lpc filter coefficients
```

```
------------------------ Allowed combinations -----------------------
                 fft      cepstre    masquage    f0    formant   lpc
partie_reelle    OK,         OK,        NO,        NO,     NO,      OK
partie_imagin    OK,         OK,        NO,        NO,     NO,      OK
frequence_hz     OK,         OK,        OK,        OK,     OK,      OK
frequ_cents      OK,         OK,        OK,        OK,     OK,      OK
frequ_midicents  OK,         OK,        OK,        OK,     OK,      OK
frequ_inst       OK,         OK,        NO,        NO,     NO,      OK
amplitude_lin    OK,         OK,        OK,        OK,     OK,      OK
amplitude_dB     OK,         OK,        OK,        OK,     OK,      OK
phase            OK,         OK,        NO,        NO,     NO,      OK
sndpressexcess   NO,         NO,        OK,        NO,     NO,      NO
truepitch        NO,         NO,        OK,        NO,     NO,      NO
weight           NO,         NO,        OK,        NO,     NO,      NO
score f0         NO,         NO,        NO,        OK,     NO,      NO
coef pitch       NO,         NO,        NO,        OK,     NO,      NO
largeur formant  NO,         NO,        NO,        NO,     OK,      NO
coeff filtre LPC NO,         NO,        NO,        OK,     NO,      OK
```

33.2 SuperVP -hf

```
===================================================
SuperVP (IRCAM) 1990-2005       version : 2.59j fix-1 (compiled by niels for
AudioSculpt  on Wed Jan 26 17:45:54 GMT 2005) Prec=double
===================================================

Filter module description ::
============================

bande,bande-noex,band-transp,band-noex-transp :
        applies a band pass/band stop filter with possibly time varying
        band edges. The parameter file contains multiple lines of

        Example: time    num  bandtype  f0 f1 f2 .... fnum

        time positions description in time
        num  number of band edge frequencies to follow
             each pair of frequencies constitutes a band
             so num has to be even
        bandtype  bands given are pass=1 or stop=0 bands
        fi   num frequencies that specify the edges of
             the alternating pass/stop bands
             included into the band are those freq. bins that
             are min(edge1,edge2) <= freq <=  max(edge1,edge2)

             The edge frequencies are grouped into bands regarding their
             position in the file. Consequently the bands may self-intersect.
             The lower edge of the band is always min(edge1,edge2)
             while the upper edge is max(edge1,edge2).
```

```
        For time varying filters only the fi entries may change with time.

        Extrapolation:

        In normal operation the first and last lines of the parameter file
        are extrapolated without changes to start/end of sound.
        If the flag -noex is appended to the filter type
        the filter will not be extrapolated but only applied inside the time
        limits specified in the parameter file

        Transparency:

        If -transp is added to the filter type the filter behavior
        is changed with respect to the mode
        of superposition CombineMax. In this case pass bands
        are considered transparent with respect to filter superposition
        such that they are not included in selecting the maximum
        value of the transfer function.

--------------------------------------------------------

breakpt breakpt-noex breakpt-transp breakpt-noex-transp:

    These filters applies a piece wise linear description
    of a time varying or constant frequency response specified in dB/rad.

    The parameter may contain an ASCII or binary description of
    the filter. In the binary case the file has to start with
    the 8 characters BPBINARY. The parameter file consists
    of sets of numbers describing the filter as a break point
    function for a given time, as follows:

    time mode numpairs freq value freq value

    time(float): time position to apply following filter response definition

    mode(int): filter mode (0=amplitude/1=phase/2=amplitude and phase)
               mode parameter can not change within a single file

    numpairs(int):  number of freq-value breakpoints to follow in current line

    freq value(float): depending of the selected mode pair or triple of values
                describing a break point of the frequency response

                freq [Hz] and amplitude [dB] for mode = 0
                freq [Hz] and phase [rad] for mode = 1
                freq [Hz], amplitude [dB] and phase [rad] for mode = 2

    The type given in parenthesis is the data type of the entries
    for the binary parameter file.

    The break point description is extended over the full frequency
```

```
    range using a constant extrapolation of the start and end break points.
    The envelopes amplitude and phase values for each frequency bin
    are linearly interpolated over time and extended to the start and
    end of the file using constant extrapolation. For correct interpolation
    the phase values are required to be unwrapped! For instantaneous
    changes with respect to time or frequency you may give the same
    frequency  or time value   twice.

    Example for breakpt file using amplitude mode :

      0.1 0 4 100 0 500 -30 1000 -30 4000 0
      0.5 0 5 100 0 500 -30 1000 -30 4000 -10 5000 0

    This file would attenuate the amplitudes in the region
    between 100 and 4000Hz to 100 and 5000Hz respectively.

    Extrapolation:

    In normal operation the first and last lines of the parameter file
    are extrapolated without changes to start/end of sound.
    If the flag -noex is appended to the filter type
    the filter will not be extrapolated but only applied inside the time
    limits specified in the parameter file

    Transparency:

    If -transp is added to the filter type the filter behavior
    and parameter file is changed with respect to the mode
    of superposition CombineMax. In this case the first floating point
    number in the parameter file is the transparent value of the filter
    with respect  to filter superposition. Filter transfer function values
    that are equal to this value are not included in selecting the maximum
    value of the transfer function.

--------------------------------------------------------

clip, clip-norm: apply clipping filter within time slices of the sound.
        The clipping filter clips all amplitudes of a spectrum
        above the upper threshold and rescales the amplitudes
        between lower threshold and upper threshold to lie between
        0 and upper threshold
        A clipping filter can be described in two modes:
        First mode describes a clipping filter operating
        on the whole sound file:

        Example: low_threshold high_threshold

        this will result in applying clipping with the given parameters
        globally for the whole sound or

        Example: low_threshold high_threshold start_time end_time
```

```
        this will apply the respective clipping parameters only within
        the specified time region. Multiple lines may be specified.
        For overlapping regions the first region takes precedence
        thresholds parameters are assumed to be in dB with respect to a
        normalized amplitude spectrum (max. = 0dB)

        Normalization:

        if the special type clip-norm is specified the clipping filter
        rescales the resulting signal to preserve the signal energy.

--------------------------------------------------------

denoise: clipping all amplitudes inside a time frequency
        surface that are below the specified clipping value
        to zero.

        The parameter file is equivalent to the surface filter with
        the gain parameter specifying a clipping level in dB relative
        to the maximum of a normalized spectrum (0dB).

--------------------------------------------------------

fof fof-noex fifof fifof-noex: filtering applies a collection
       of second order resonance or formant filters to the sound signal.

    The file consists of lines containing

    time numfof freq gain band ...

    time:    start time for the following parameters
    numfof:  number of formats to follow

    then numfof format descriptions each specifying

    freq:   fof resonance freq in Hz
    gain:   amplitude gain at resonance freq in dB
    bandwidth: formant bandwidth in Hz


    In case of fof the number of formant filters has to be the same
    throughout the file and the parameters for each formant are linearly
    interpolated.
    In case of fifof the number of formants can change and interpolation
    with respect to time is achieved by interpolating the resulting
    transfer functions.

    example:    fifof.par

    0.1 1 440 0 50
    0.2 2 440 0 50 880 -2 100
    0.3 3 440 0 50 880 -2 100 1760 -4 200
```

```
    Extrapolation:

    In normal operation the first and last lines of the parameter file
    are extrapolated without changes to start/end of sound.
    If the flag -noex is appended to the filter type
    the filter will not be extrapolated but only applied inside the time
    limits specified in the parameter file

--------------------------------------------------------

fshift:  this filter applies a time varying
        frequency shift to the signal

        the parameter file contains a breakpoint
        description of the shift to apply at a given time
        point in Hertz. The first and last parameter
        are extended till the start/end of the sound respectively

        example:

        0.1 -50.

        this file would shift the sound down by 50Hz.

        NOTE:    due to the internal representation you may not
        combine a frequency shift filter with time stretching
        or transposition.

--------------------------------------------------------

gabarit: apply a fft analysis output file that has been
        stored in gabarit format as a filter to the input sound file.

    The gabarit mode has to be 1 (amplitude only), 2 (phase only) or 3
    for amplitude and phase filtering (See option -Og). Since SuperVP 1.83
    the filter analysis data sampling rate and fftsize are no longer
    required to match the respective values for
    the filtered sound. Any mismatch will result in a
    proper resampling of the filter data using linear interpolation
    of phase and amplitude (amplitude interpolated in log domain) such
    that the frequency response is kept.

    Gabarit format description

      All the different formats share a common file header specifying

        gabarit mode (4 byte int)
        sample rate (4 byte float)
        soundfilenamelen (4 byte int)
        soundfilenamechars (namelen chars)
```

```
        The length of the input sound file is always including
        the limiting zero byte of a c-string and rounded
        up to regard a 4 byte boundary, the unused part of the name
        is filled with zeros.

        Following the file header are the frames
        each with a unified frame header consisting of

        frameCenterTime (4 byte float)
        nbdata          (4 byte int)
        windowSize      (4 byte int)
        frameSize       (4 byte int)

        The samplerate is used to derive the data spacing in frequency as
        samplerate/(nbdata-1)/2. Windowsize is used to adjust the phase
        prior to using it by adding/subtracting a linear phase component that
        compensates for the center of the window being shifted (windowsize-1)/2
        away from the start of the frame. Framesize is not used
        for filtering and can be left 0. After each header there are
        nbdata multi element values. The meaning of the values
        depend on the mode and are:

        mode  |  data
        ==============================
         1    |  single amplitude
         2    |  single phase value
         3    |  pair of amplitude and phase value
         4    |  pair of frequency and amplitude value (not suitable for filtering)
         5    |  pair of real imaginary value (not suitable for filtering)

         All data points are in 32 bit float format

      The phase interpolation is done after phase unwrapping. Therefore,
      the maximal spacing of synthetic phase transfer functions has to be
      selected such that phase differences of the phase transfer
      function between the bins are smaller then PI.
--------------------------------------------------------

surface : Applies constant gain to a piece wise
          linear surface of the spectrogram.

          The parameter file contains multiple lines for each surface,
          and may contain multiple surfaces. General syntax

          First line of a surface:
             start_time num end_time gain_dB

             start_time/end_time:  the minimum and maximum
                  time covered by that surface.

             num: the number of lines to follow describing the shape of
                  the surface.
```

```
           gain_dB: The fixed attenuation/amplification for this surface.

        surface description is given by num lines of format
           time low_freq high_freq

           time: specifies the time the description applies
           low_freq/high_freq: specify the frequency boundaries
               (end points belonging to the surface) for the given time.

        The description specifies a piece wise linear bounded surface
        in the time/frequency plane. The start_time/end_time locations
        of the surface  are centered between the first and last
        frequency boundary pairs.
        Due to this format only convex surfaces can be described.

        Example for single surface with  40dB attenuation:
        0.0 4 5.0 -40
        1.0 1000 2000
        2.0 1000 2000
        3.0 2000 3000
        4.0 2000 3000

        Time and frequency values have to be non decreasing.
        Superposition of different surfaces will respect the
        selected superposition mode.
--------------------------------------------------------

sub, sub-det, sub-noex, sub-noex-det:
        spectral subtraction filter will subtract
        time interpolated noise spectrum that has been created via -OM1/-OM2
        output option from noise segments specified via -avseg option
        from the signal spectrum. The spectral subtraction allows fine control
        via additional command line parameters: -avsfac, -avbeta,
        -avgamma, -avrelax.

        The parameter file contains the spectral description of
        the noise spectra for all the segments. According to the selected
        description mode the first column contains the mean or maximum value
        M(k) for each bin k and the current segment. The second column contains
        the standard deviation S(k) of the bin for the current segment.
        The noise estimate for bin k is given by

        NE(k) = M(k) + 10^(avsfac/20) * S(k)

        If we denote NI(k) the interpolation of the two noise
        estimates from the neighboring segments contained
        in the parameter file the spectral
        subtraction module would create the output amplitude
        spectrum O(k) from the input amplitude I(k) spectrum following

        O(k) = max(0,Gl(k)*I(k)) with the generalized spectral subtraction factor
```

```
        G(k)  = max(sum_bark([1-[NI(k)/I(k)]^avgamma]^(1/avgamma)),10^(avbeta/20))
        Gl(k) = max(Gl(k) * exp(-td/(avrelax/1000)),G(k))

        G(k) will  be close to 1 if the input amp. spectrum is larger
        than the noise estimate and become small if the input amp.
        spectrum approaches the noise estimate. Gl provides a means
        to modify the decay rate of the spectral subtraction factor.

        To prevent musical noise the spectral subtraction factor G(k)
        is always averaged over the bark band related to bin k.

        Meaning of command line parameters:
        ----------------------------------

        avbeta   : controls remaining noise floor (maximum spectral attenuation
             is given by avbeta).
             see: control flag -avbeta
        avgamma : switches between amplitude (gamma=1)
             and energy (gamma=2) subtraction.
             see: control flag -avgamma
        avsfac   : controls impact of standard deviation on the noise estimate
             and establishes a means to obtain over/under-compensation.
             see: control flag -avsfac
        avrelax : do not decrease Gl(k) faster than with an
             relaxation time constant. avrelax is given in milliseconds.
             see: control flag -avrelax

        Deterministic noise:
        ====================

        By adding the flag -det to the filter type special treatment of
        stationary sinusoids in the noise estimate is requested.
        Sinusoids are detected from the averaged spectrum
        and removed below the noise level prior to applying
        spectral subtraction.


        Extrapolation:

        In normal operation the first and last lines of the parameter file
        are extrapolated without changes to start/end of sound.
        If the flag -noex is appended to the filter type
        the filter will not be extrapolated but only applied inside the time
        limits specified in the parameter file

--------------------------------------------------------

-avbeta <noise floor value> :
        Specifies the noise floor to keep when applying the spectral
        subtraction filter. The value is specified in dB. (Def: -100dB)
```

```
        name is given this file should contain pairs of start and end times
        of segments that will be used to calculate averaged spectra.
        If the flag is followed by two numbers they are used as start
        and end times of the next segment to be added to the list of
        segments. The parameter can be specified multiple times,
        however, the segment positions need to be in increasing
        time order.

        ---- Example: -avbeta -20

        Limits noise attenuation in the spectral subtraction filter
        to 20dB.

        See: sub filter module

----------------------------------------------------------

-avsfac <standard deviation multiplicand> :
        Specifies the factor to multiply with the noise standard deviation
        for estimating the noise amplitude spectrum in the
        spectral subtraction filter (Def: 0).

        ---- Example: -avsfac 1.

        Adds one time the standard deviation to the mean amplitude
        to obtain the noise estimate.

        See: sub filter module

----------------------------------------------------------

-avgamma <1 or 2> :
        Specifies amplitude or energy subtraction mode for the
        spectral subtraction filter (Def: 1).

        ---- Example: -avgamma 2

        Selects energy control instead for amplitude control
        for the spectral subtraction filter.

        See: sub filter module

----------------------------------------------------------

-avrelax <relaxation time constant> :
        specifies the time constant in ms for the decay
        of the generalized filter factor in the spectral subtraction
        filter (Def: 0.).


        ---- Example: -avrelax 20.
```

```
Relaxation time constant of generalized subtraction filter
set to 20ms.

See: sub filter module
```
